DOCUMENT RESUME

ED 275 740 TM 860 654

AUTHOR Resnick, Lauren B.
TITLE Cognition and Instruction: Recent Theories of Human Competence and How It Is Acquired.
INSTITUTION Pittsburgh Univ., Pa. Learning Research and Development Center.
SPONS AGENCY National Inst. of Education (ED), Washington, DC.
REPORT NO LRDC-1986/6
PUB DATE Jun 86
NOTE 66p.
PUB TYPE Information Analyses (070)

EDRS PRICE MF01/PC03 Plus Postage.
DESCRIPTORS Association (Psychology); *Cognitive Ability; *Cognitive Processes; *Cognitive Psychology; *Competence; Educational Research; *Epistemology; Instruction; Intelligence; Language Processing; Learning Processes; Mathematics Skills; Metacognition; Problem Solving; Reading Processes; Research Methodology; Schemata (Cognition); *State of the Art Reviews; Teaching Methods

ABSTRACT

This paper outlines the current state of knowledge about how intellectual competence is acquired. Implications resulting from recent changes in the psychology of learning and development are beginning to revitalize the science of learning and instruction. Four broad topics illustrate the links between instructional experiments, fundamental research on learning and thinking, and potential applications to the science of instruction: (1) understanding natural language; (2) learning to read; (3) developing mathematical competence; and (4) problem solving, intelligence, and learning abilities. Major themes emerging from research in language processing indicate that prior knowledge is essential in constructing meaning for a new text and that the construction of meaning centrally involves inference. Two main themes in reading research are emphasized: (1) the active interplay between expectations and the visual stimuli of printed words; and (2) the central role of automatic processes of word recognition. Special problems arise when analyzing mathematics as a domain of cognition and learning, including organizing schemata, recognizing persistent and systematic errors, and linking symbols and their referents. Recent work in problem solving is focused on performance in information-rich domains. There has been considerable effort in recent years to reanalyze the constructs of intelligence and aptitudes in terms of cognitive processes and constructs, including metacognitive skills. Directions for further research are suggested. (LMO)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

ED275740
TM 860654

LEARNING RESEARCH AND DEVELOPMENT CENTER

1986/6

COGNITION AND INSTRUCTION: RECENT THEORIES OF HUMAN COMPETENCE AND HOW IT IS ACQUIRED

U.S. DEPARTMENT OF EDUCATION
Office of Educational Research and Improvement
EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)

This document has been reproduced as received from the person or organization originating it.

☐ Minor changes have been made to improve reproduction quality.

• Points of view or opinions stated in this document do not necessarily represent official OERI position or policy.

"PERMISSION TO REPRODUCE THIS MATERIAL HAS BEEN GRANTED BY

J. Ang

TO THE EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)."

University of Pittsburgh

COGNITION AND INSTRUCTION: RECENT THEORIES OF HUMAN COMPETENCE AND HOW IT IS ACQUIRED

Lauren B. Resnick

Learning Research and Development Center
University of Pittsburgh

1986

Reprinted with permission from B. L. Hammonds (Ed.) Psychology and learning: The master lecture series, Vol. 4 (pp. 123-187). Copyright 1985 by American Psychological Association.

Preparation of this chapter was supported by the Learning Research and Development Center, funded in part by the National Institute of Education (NIE), U. S. Department of Education. The opinions expressed do not necessarily reflect the position or policy of NIE, and no official endorsement should be inferred.

LAUREN B. RESNICK

COGNITION AND INSTRUCTION: RECENT THEORIES OF HUMAN COMPETENCE

4

Following receipt of her undergraduate degree from Radcliffe College, Lauren B. Resnick earned a master's in Teaching and an EdD in Research in Instruction from the Harvard Graduate School of Education. Resnick is currently a Director of the Learning Research and Development Center and Professor of Psychology and Education at the University of Pittsburgh. She is past president of the Division of Educational Psychology of the American Psychological Association (APA), and a fellow in the APA's Divisions of Experimental Psychology, Developmental Psychology, Educational Psychology, and the Experimental Analysis of Behavior. She also has served as vice president of the American Educational Research Association's Division of Learning and Instruction, is a fellow of the American Association for the Advancement of Science, and is a member of the National Academy of Education.

Resnick's primary interest is the emerging field of the cognitive psychology of instruction, and she is playing an active role in its development. Her major current research focuses on the learning of mathematics and science. She has also studied reading and other subject-matter domains. In addition to her research efforts, Resnick is founder and editor of *Cognition and Instruction,* the major new journal in the field.

5

LAUREN B. RESNICK

COGNITION AND INSTRUCTION: RECENT THEORIES OF HUMAN COMPETENCE

Introduction

My goal in this essay is to sketch the current state of knowledge about how intellectual competence is acquired and to suggest directions for future research, especially research that promises to improve instruction. My task is more complex and more exciting than it would have been 10 years ago because the psychology of learning and development has in that time undergone a profound change. In the past decade, a number of the assumptions that had guided research on learning have been called into question, and a vigorous new "cognitive science" has taken hold. The implications of this change for conceptions of human mental functioning are vast and the possibilities for a revitalized science of learning and instruction are just beginning to be realized.

Until 10 or 15 years ago, psychologists interested in the nature and acquisition of intellectual competence were faced with some unpalatable choices. One set of psychological theories—those in the associa-

The research reported herein was supported by the Learning Research and Development Center, funded in part as a research and development center by the National Institute of Education (NIE), Department of Education. The opinions expressed do not necessarily reflect the position or policy of NIE, and no official endorsement should be inferred.

tionist and behaviorist traditions—offered a rigorous experimental approach and an active concern not only for the nature of changes in performance but also for the means to influence change. These theories provided a strong basis for studying and prescribing interventions and were attractive to those who wanted to have an impact on instruction and education. But the associationists and behaviorists had little to say about the nature of thought processes or the questions of structure, organization, and meaning in learning. In their insistence that overt behavior was the only proper object of scientific study and their attempt to reduce thought to collections of associations, these psychologists offered no theoretically sensible way of dealing with questions of understanding and provided little wisdom concerning the nature of intellectual competence or the role of structure and meaning in learning.

On the other hand, structuralist theories (those of Piaget and the Gestalt psychologists, for example), which did treat mental life as real and important, and which offered strong theories about the role of structured knowledge and meaning in intellectual competence, had very weak theories of acquisition. And they had even less guidance to offer on how to intervene in acquisition. Despite elegant examples of the kinds of instructional goals that might be promoted, neither Piagetian nor Gestaltist analyses proceeded very far in specifying goals for instruction in anything like the rigorous detail that learning theorists offered for the associations or behaviors to be fostered by instruction. Furthermore, the relative silence, at least for Piagetians, on questions of instruction was further heightened by a kind of mistrust of intervention that was inherent in biologically oriented developmental theories (cf. Resnick, 1981a). Psychologists thus were faced with a choice between (a) theories that were centrally concerned with changes and how to promote them, but fundamentally unconcerned with thinking and meaning and (b) theories that were centrally concerned with structures of thinking and with understanding, but very vague about mechanisms of acquisition and disinterested in or even mistrustful of instruction.

Recent developments in cognitive psychology offer a new set of perspectives on this troubling dichotomy. The heart of cognitive psychology is the centrality given to the human mind and the treatment of thinking processes as concrete phenomena that can be studied scientifically. Researchers in various branches of psychology have found common ground in the study of cognition, and they have been joined by computer scientists, linguists, and philosophers to form a new cognitive science research community. This new interest in cognition has resulted in (a) a flourishing of research on complex forms of knowledge and skill, (b) a convergence on some key points between experimental and structuralist tradition in psychology, and (c) the development of a variety of new methods of research and forms of theorizing that are gradually developing a new scientific method specifically suited to the study of human mental functioning.

One of the most important developments in cognitive science has been the gradual construction of new ways of linking knowledge and performance. Process theories of cognitive functioning provide precise statements of how the knowledge that people possess permits them to perform in certain ways on certain kinds of tasks. The interest in processes of thought has led to the refinement of methods that trace sequential steps in thinking. These methods include recording patterns of reaction times for stimuli or tasks of different complexity, tracking eye movements as subjects read texts or solve visually presented problems, and using "think-aloud" protocols in which subjects solve problems while verbalizing what is going through their minds as they work. Because think-aloud methods seem to share features of the long-discredited introspective methods of psychology, they have evoked a certain degree of skepticism. Careful methodological work (e.g., Ericsson & Simon, 1984) has established the limits and powers of these methods.

Study of the relations between processes and content of thought is further stimulated and strengthened by the active enagagement of psychologists with computer scientists, especially those interested in the study and development of artificial intelligence. Viewing the computer as a metaphor for the human mind has stimulated cognitive psychology, allowing for more intentional and goal-driven processes than the older image of the mind as a switchboard, which was so neatly compatible with associationist theories. However, the real power has come not from the general metaphor, but rather from the use of computer programs as detailed simulations of human thinking and, thus, as a way of both energizing and disciplining cognitive theory. When computer programs behave as humans do—making similar mistakes, pausing at similar points, expressing confusion over the same issues—it is reasonable to assume that the internal processes of the human and the computer are similar, and researchers can treat the programs' visible processes as a theory of the invisible processes of humans.

Initially applied to limited forms of problem-solving task performance (Newell & Simon, 1972), computer simulation as a form of psychological theory has subsequently been extended to a wide variety of tasks and domains, and more recent work provides complex models of how knowledge is structured and accessed in addition to the procedures and heuristics used in manipulating it. This use of computer programs as models of human thinking has been enhanced by important shifts within artificial intelligence itself (Dehn & Schank, 1982). Artificial intelligence researchers, finding that truly complex forms of thinking depend on optimally structured knowledge and heuristic rather than exhaustive forms of searching this knowledge, are turning more and more to studies of human intelligence to inform their efforts to build intelligent machines. In this emerging field of cognitive science, it is not always easy to tell who is a psychologist and who is an artificial intelligence specialist.

8

In addition to having an interest in thinking and recognizing the central role of structured knowledge in the process of thought, the modern cognitive science converges with structuralist traditions in psychology in rejecting the long-held *tabula rasa* assumptions of American psychology in favor of what can best be called a *constructivist* assumption. In associationism there was no way to imagine knowledge entering a human's mind except from the outside. Objects could be perceived, associations between events noted, and mental bonds gradually built. These bonds were, it was assumed although often not directly discussed, direct reflections of the external information to which one was exposed. To learn was to build up more and more of these records and to make them more quickly accessible. But to learn was not to construct new associations and relationships through purely mental activity.

Today's cognitive science, by contrast, gives a central place to organizing structures and thus provides the terms in which theories of how individuals build new relationships can be developed. Knowledge is no longer viewed as a reflection of what has been given from the outside; it is a personal construction in which the individual imposes meaning by relating bits of knowledge and experience to some organizing schemata. This constructivist view in cognitive science is not identical to Piagetian constructivism, but it is close enough in spirit that psychologists who a decade ago could find little ground for serious debate can now successfully respond to and build upon each others' work. One result has been a rejoining of forces by certain groups of developmental and experimental psychologists who had for some decades diverged in their interests.

I will illustrate and elaborate all of these trends and the research methods on which they are based in the course of this chapter. I will also stress another important characteristic of recent cognitive research—that is the extent to which it is both relevant to and driven by questions concerning instruction and the deliberate modification of human competence. Partly because cognitive scientists are seeking complex and "ecologically valid" domains of human intellectual functioning in which to develop their theories, and partly because of a drive toward socially relevant applications of their work, cognitive psychologists are devoting substantial effort to research on the kinds of tasks that are studied in school or other educational institutions.

Both the nature of competence in such domains and its acquisition are increasingly central questions in today's research. Instructional experiments, when conducted so as to reveal details of the learning process, are a valuable tool in research on processes of acquisition, and these experiments further tighten the links between fundamental research on learning and thinking and potential applications to a science of instruction. To illustrate all of this, I will build my chapter around four broad topics: understanding written and spoken language; learning to read; developing mathematical competence; and the nature of problem solving, intelligence, and learning abilities. These are all do-

mains in which the convergence of basic and applied research on human cognition is highly evident and in which many of the methodological tools and theoretical issues of current cognitive science can be displayed.

Understanding Natural Language

I will focus first on how people understand that most complex of human intellectual productions, language. This question has captured the attention of some of the world's best psychologists, linguists, and computer scientists over the last 15 years, and the result of their work is a rich body of knowledge and theory about how people understand what they read or what other people tell them.

Contrary to certain older views, cognitive scientists now agree that the process of understanding language is not one of absorbing and recording what is written or said. Rather, in this process the message is used to build up a representation in one's mind of the situation to which the message refers. This representation is simultaneously selective and elaborative with respect to the message. It does not exactly match the message. Rather, some things that the message says are left out, and some other information that the message left out is put in. The mental representation is elaborated by the reader or listener to include things not stated explicitly but necessary to make sense of the message. Information that the receiver construes as not being crucial to the meaning is left out of the receiver's mental representation. This process of constructing representations based on messages highlights a central feature of natural language understanding: Except in special circumstances, it is not the message itself that is represented, but its reference. People use language to refer to something external to the language itself, and the processes of language interpretation are all aimed at understanding that external situation. Knowledge of linguistic conventions as such, while crucial to the process of understanding, is normally employed to aid understanding of the reference situation rather than as an end in itself.

The processes by which the referential meaning of a message is constructed by a reader or listener have been a central concern of cognitive scientists interested in natural language understanding. Two major themes emerge from this work. First, prior knowledge is essential in constructing meaning for a new text. Second, the construction of meaning is one that centrally involves inference.

The Role of Prior Knowledge in Constructing Representations: Schemata in Language Processing

An example from a now classic experiment in cognitive psychology is the best way to demonstrate the importance of prior knowledge in

understanding a text. Read the text in the following paragraph, but do not look ahead to Figure 1 as you do so.

> If the balloons popped the sound wouldn't be able to carry since everything would be too far away from the correct floor. A closed window would also prevent the sound from carrying, since most buildings tend to be well insulated. Since the whole operation depends upon a steady flow of electricity, a break in the middle of the wire would also cause problems. Of course, the fellow could shout, but the human voice is not loud enough to carry that far. An additional problem is that a string could break the instrument. Then there could be no accompaniment to the message. It is clear that the best situation would involve less distance. Then there would be fewer potential problems. With face to face contact, the least number of things could go wrong. (Bransford & Johnson, 1972, p. 719)

Unless readers know of the Bransford and Johnson (1972) experiment and thus remember what the text is about, virtually everyone reading this text has the experience of not understanding. The text seems garbled and senseless. Now look at Figure 1; it tells you, via a pictorial illustration, what the text is about. After seeing this serenade picture, most people experience a sense of insight concerning the text. They are ready to say, "Now I understand." The framework provided by the picture provides a "scaffolding" for interpreting the text.

The text in this study was a particularly ambiguous one, deliberately chosen to show that prior knowledge about the reference situation is crucial in understanding a text. Yet the same phenomenon has been observed in far less extreme situations as well. Hints provided in advance by the experimenter or the reader's own background have been shown to make a difference in what the reader understands in a text. For example, one study (Anderson, Reynolds, Schallert, & Goetz, 1977) shows that music students interpret the following passage as a description of an evening of playing chamber music, whereas physical education students interpret it as a story about an evening of card playing. This kind of study shows clearly that the background knowledge and interpretive schemata that readers bring with them to a text make a difference in what they understand the text to be saying.

> Every Saturday night, four good friends get together. When Jerry, Mike, and Pat arrived, Karen was sitting in her living room writing some notes. She quickly gathered the cards and stood up to greet her friends at the door. They followed her into the living-room, but as usual they couldn't agree on exactly what to play. Jerry eventually took a stand and set things up. Finally, they began to play. Karen's recorder filled the room with soft and pleasant

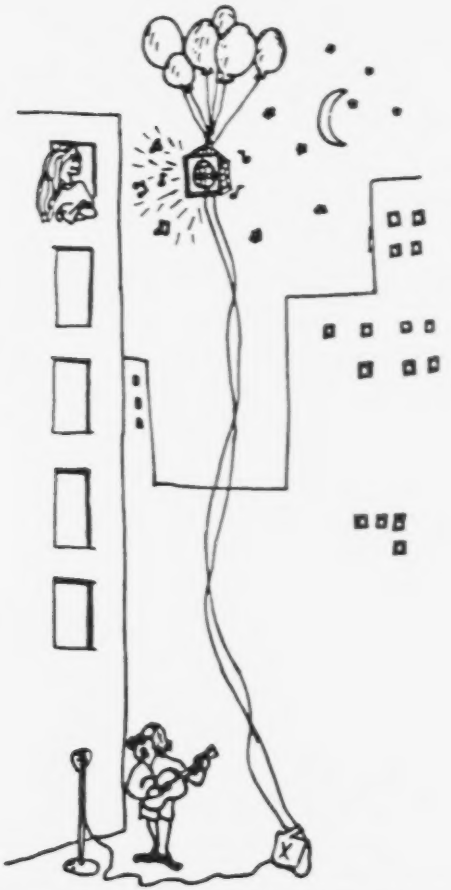

Figure 1. Appropriate context for ambiguous textual passage. Reprinted from Bransford and Johnson (1972) by permission.

music. Early in the evening, Mike noticed Pat's hand and the many diamonds. As the night progressed, the tempo of play increased. Finally, a lull in the activities occurred. Taking advantage of this, Jerry pondered the arrangement in front of him. Mike interrupted Jerry's reverie and said, "Let's hear the score." They listened carefully and commented on their performance. When the comments were all heard, exhausted but happy, Karen's friends went home. (Anderson, Reynolds, Schallert, & Goetz, 1977, p. 372)

Demonstrations of these kinds, together with many more formal experiments, underlie what has become known as the *schema theoretic* view of language comprehension. The theory holds that *schemata*, which are prototypical versions of a situation, are stored in people's minds and are used to interpret new instances and events in those situations. The schemata describe classes of situations and specify the relations between objects and events. The specific events and objects vary according to the particular case, but the relations specified in a more general schema still hold and are used to interpret the case at hand. Schemata result from a cumulation of prior learning and experience. They are necessary if one is to comprehend new verbal material, and thus they are important to all learning that depends on verbal presentations. This schema-based view of understanding goes back to early work in experimental psychology by Bartlett (1932), which showed that when a text was recalled, elements were deleted or highlighted according to a directing interpretive schema. Schemata of this kind, given various labels (such as "scripts," "frames," "memory organization packets"), are at the core of all artificial intelligence models of language understanding (see Dehn & Schank, 1982; Schank & Abelson, 1977, for discussion of these points).

In the serenade example, consider what might have happened to produce the initial failure to understand and the subsequent sense of understanding without difficulty. It is not the words of the text itself that produce the difficulty; English language readers can attach meaning to each word, and every sentence is grammatically correct. In fact, read by itself, each sentence is understandable. The problem occurs when the reader tries to make the sentences fit together in a coherent whole. Pairs of adjacent sentences seem to have no connection to one another. Once the context of the serenade is known, however, the reader can infer the connections and the passage makes sense.

With this simple analysis, I have already identified several elements of the process of understanding: Readers and listeners must access previously stored knowledge about the meaning of individual words; they must use their knowledge of syntactic rules and conventions and of the world to sensibly convert phrases and sentences in the text into propositions about a situation; and they must link the propositions into a coherent representation of a single situation. An excellent description of the various kinds of processing activity used to comprehend a text is given by Perfetti (in press). Perfetti shows that even in assigning meaning to individual words and analyzing sentences into sensible propositions about the world (both processes that proceed largely automatically without conscious effort or attention), prior knowledge about the reference situation and the conventions of language play a powerful role. After the propositions are developed, other processes and knowledge are used to link them into a coherent representation of the reference situation for a text.

Inference and Coherence-Building in Understanding Language

All natural language communications are incomplete because they do not specify everything about the reference situation that is needed for a complete and coherent representation. To build a coherent representation, readers or listeners must use their knowledge to infer links between individual propositions in a text and to provide a framework in which to interpret specific information supplied in the text. The work of Kintsch and vanDijk (1978) and their colleagues is the most extensively developed theory of the process of building coherence. I use an example based on Kintsch (1979) to illustrate.

The brief text passage that follows this paragraph is broken down into numerous propositions. These propositions are the elementary pieces of information conveyed. Sentences may contain one or more propositions. A text is said to be locally coherent to the extent that each new proposition makes explicit reference to recently stated prior propositions. Proposition sequences 1–4 and 5–11 are fully coherent because the actor in each proposition has already been named. This means that these segments are coherent within themselves. However, the two parts are not coherent if the sequences are joined, because proposition 5 is not explicitly linked to its predecessors. To understand this text the reader must infer a proposition that will link proposition 5 to its predecessors. Such a proposition might be: "The Swazi tribe had warriors."

Text

The Swazi tribe was at war with a neighboring tribe because of a dispute over some cattle. Among the warriors were two unmarried men named Kakra and his younger brother Gum. Kakra was killed in battle.

Propositional analysis of text

1. The Swazi tribe was at war.
2. The war was with a neighboring tribe.
3. The war had a cause.
4. The cause was a dispute over some cattle.
5. There were warriors.
6. The warriors were two men.

7. The men were unmarried.

8. The men were named Kakra and Gum.

9. Gum was the younger brother of Kakra.

10. Kakra was killed.

11. The killing was in a battle. (see Kintsch, 1979)

In extensive research on the processes of building local coherence, Kintsch and vanDijk and their colleagues have shown that features of the text, such as the number of missing propositions and the distance in the text that must be traversed to find explicit links, affect how long it takes to read the text and how easy it is to understand the text. Kintsch and vanDijk's theory accounts for a large body of such findings. This theory assumes—as do other modern information processing models—that human capacity for holding information in *immediate* (working) memory and for operating on this information is limited. For this reason, one cannot imagine that in the process of building a representation of a text, all propositions that have been read are brought into working memory every time a search for coherence is made. Instead, the theory assumes that reading and representation-building occur in cycles, where each cycle represents an attempt to link one or more new (just read) propositions to the representation already built. Because of working memory limitations, not all of the representation can be held in working memory—and therefore be searched—on a given cycle.

The ease or difficulty of comprehending and the time that comprehension takes depends, therefore, on whether the particular part of the representation retained in working memory or in a given cycle contains a reference to which the next proposition can be linked. If not, a new choice of propositions will have to be made, which produces time delays and at least temporary hesitation and confusion. Immediately preceding propositions from the text are always likely to be in working memory, and this accounts in part for the fact that when links can be created between adjacent propositions, comprehension proceeds more smoothly than when links must be created with propositions stated some time earlier. For more distant (earlier) propositions, ease of comprehension will depend on whether the choice of propositions to retain in working memory has been felicitous. This depends, in turn, on the extent to which the text provides clues as to what information is most important and the extent to which the reader or listener is adept at using these clues. This brings me to the important question of how people know, and how texts signal, what is most important in a verbal message.

Macrostructures and Frameworks for Interpretation

For most written texts, readers tend to agree fairly well on which statements are important or central and which ones are subordinate, perhaps functioning only as elaborations. Meyer (1975) has used this regularity to develop a method of coding the statements in a passage for their relative centrality; other less systematic ways of judging centrality have also been developed. Using measures of this kind, it has been possible to show that the material most likely to be forgotten or left out of a summary is the material lowest in the hierarchy of importance. Conversely, if material high in the hierarchy is not specified in the text, people will have trouble interpreting the text at all, will tend to insert missing high-level propositions in their summaries, and will spend a long time studying the portion of the passage where the high-level organizing material is expected to be (Kieras, 1977). Also, when asked whether a given statement was or was not present in the text, people are likely to assert with great confidence that highly central material that is consonant with the main theme of the text was there—even when it was not.

Voss and his colleagues (Chiesi, Spilich, & Voss, 1979; Spilich, Vesonder, Chiesi, & Voss, 1979) have shown that readers' ability to make inferences depends upon what they already know about the topic of the text. Finally, several studies have shown that more competent readers are better able than weaker readers to detect the hierarchy of importance in a text, picking out not only the main idea but also layers of supporting argument and detail (Meyer, 1984). I have already shown that knowing what a text is about plays a role in understanding it. However, in ordinary reading—unlike the serenade example—the necessary information about the theme of a text is not provided externally by a picture, but rather must itself be inferred in the course of reading.

In the process of successful comprehension, readers not only build up local coherences between propositions but also develop a representation of the *gist* of the message. A gist representation includes only the most important information given in the text; details are dropped out. But the gist representation is not just a string of important individual elements. It is itself organized so that these elements make sense with respect to one another. Kintsch and vanDijk have called these gist structures "macrostructures," to distinguish them from the "microstructures" that are constructed as individual propositions that are related to one another. They have elaborated a theory of how macrostructures are created by the reader, who uses special operators to pick out and combine elements of microstructure. This can successfully be done only when appropriate schemata already in the readers' (or listeners') long-term memories are found and applied.

Macrostructure representations are built up gradually in the course of reading or listening. If the initial macrostructure provides a sensible

framework for the entire text, then the macrostructure will be elaborated and refined but not essentially modified in the course of understanding. Sometimes, however, the framework that has been guiding text interpretation turns out to be inappropriate as more of the text is read or heard. The tribal war story allows me to illustrate this point and also to emphasize again the central and multiple roles that prior knowledge plays in understanding any verbal communication. For the part of the story analyzed thus far (propositions 1–11), the macrostructure that is constructed concerns tribal war. The initial sentence would surely evoke such a theme and subsequent sentences do not disturb such an interpretation. However, the actual text from which the excerpt is drawn in fact goes off in another direction. The next sentence is: "According to tribal custom, Kakra was married subsequently to the woman Ami." A propositional breakdown of this sentence would include the propositions shown in the following paragraph.

Propositional analysis of text

12. Kakra was married.
13. The marriage was after Kakra was killed.
14. The marriage was to a woman.
15. The woman was named Ami.
16. The marriage was in accord with tribal custom. (see Kintsch, 1979)

It is easy to find the links that make these propositions coherent at a microlevel. It is not even difficult, if only microlevel coherence were in question, to link it to the preceding representation. Only a couple of propositions back there is reference to the actor in the first new proposition, Kakra. This, however, cannot be a full model of how humans understand a text, for all readers immediately recognize an anomaly and refuse the simple linkage at the propositional level. How can Kakra, who has been killed, now be married? Theories of understanding must be able to explain how such anomalies are recognized. All such theories base the recognition on prior schema-like knowledge: People have a schema for killing that produces an automatic inference that Kakra is dead; people have a schema for marriage that requires that husbands be alive. This means that Kakra cannot fill the *husband* slot in the marriage schema, so a fully coherent representation of the text cannot be built without further information.

There is in fact another schema that, if available, would solve the problem. The schema concerns ghost marriage, a tribal custom in which the oldest son of a family who dies without heirs is subsequently married with his younger brother taking his place until an heir is produced. A ghost marriage schema would provide a slot, not for a live husband,

but for a dead man and a younger brother. Kakra can fill the dead man's slot, Gum the younger brother's, and now a representation coherent at a macrolevel of interpretation can be constructed. This new representation will also contain a new candidate for the organizing theme of the story: It now appears that a summary of the story ought to include a ghost marriage, and it might well be that as the text continued, the theme of tribal war would disappear at the macrostructure level in favor of the custom of ghost marriage.

I have focused here on a particular example and on one theory of how text representations are built. Many other investigators have explored how high-level organizing information—that is, macrostructures—controls and supports the process of comprehending a text. A particularly well-developed domain for this research has been story understanding. Several investigations (see Stein & Trabasso, 1983) have shown that there is a prototypical structure of narratives that is used by people to interpret stories. The idealized story, in effect a schema of a story, organizes and directs peoples' interaction with the particular story they are reading or hearing. The story schema specifies the types of information that should be presented and the types of logical relationships that should link the story elements. Several categories of information must occur in order: a setting, an initiating event, an internal response, an attempt to obtain a goal, an outcome or consequence, and a reaction.

Some of the categories in this structure are more central than others. This is shown by substantial regularities in the portions of stories that people omit and the portions they add when asked to retell stories they have heard. Initiating events, attempts to achieve a goal, and consequences are nearly always included, but other categories, especially internal cognitive responses of the characters, are likely to be omitted (Mandler, 1978; Mandler & Johnson, 1977; Stein & Glenn, 1979; Thorndyke, 1977). Story comprehension and recall are also sensitive to the order in which categories of information are presented. People have difficulty recalling stories when information is given in an order other than that specified in the idealized story schema, and they tend to recall story information in the order predicted by the schema even when the text from which they learn the story uses a nonstandard order.

Recent research (see Flammer & Kintsch, 1982) suggests that the semantic content, rather than just form or placement of the information within the story, may be determining recall. Attempts to enlarge research on story understanding beyond the simple demonstration of story schemata (sometimes called "story grammars") have been leading psychologists increasingly to study the specific kinds of social knowledge held by children of different ages and stages of development. The newest research on story understanding suggests that widely shared knowledge about goals, plans, actions, outcomes, and motives (e.g.,

Voss, 1984) is at the heart of story understanding. Trabasso, Secco, and van den Broek (1984), for example, present evidence that knowledge about physical and psychological causality plays a central role in understanding stories. The relative weights and possible interactions between this kind of general knowledge of the world and knowledge of rhetorical structures such as story grammars is a topic of much debate in the field today.

One of the features of the Kintsch and vanDijk theory just outlined is that the processing of texts is assumed to be more or less sequential; that is, people build up their representation of the reference situation of the text bit by bit, as they go along. This means that the process of interpretation is continuous. People do not hold pieces of uninterpreted text in mind for a period of time and then later reflect on its meaning. Another important line of research on reading that has used quite different methods of study confirms this sequentiality. Just and Carpenter have for a number of years been studying reading, using eye-movement records as their basic data. They have constructed a model (a computer simulation program) of the reading process (Thibadeau, Just, & Carpenter, in press) that accounts for the patterns of eye movements observed in subjects while reading. This model, READER, processes the text in a largely word-by-word fashion. As it encounters each new word, it finds the meaning of the word and more or less simultaneously uses schemata and related semantic processing mechanisms to build up a representation to that point. It does not, in other words, delay interpretation until a whole phrase or sentence has been read. Furthermore, it builds its representation using a combination of expectations for what *should* appear next based on the context and the information actually in the printed text.

There is a striking degree of convergence between the different kinds of available evidence for how people understand written texts. It appears first that what is done automatically can also be done consciously; that is, some portions of think-aloud protocols produce sequences of steps that are not very different from those of automatic processing, in which people make successive links between sentences and store up partial interpretations as they go. However, when difficulties are encountered in the course of reading, skilled readers seem to use conscious processes to resolve the problem. In the time-course studies of reading, these are the points at which very long delays occur, and the protocol analysis studies provide a good sense of what is happening at those points of delay. At these times, there is a considerable amount of looking back, of reconstructing, and of forming or accessing new schemas for interpretation. Thus, the studies that focus on automatic processing and those that focus on conscious processing reveal similarities that seem to create a plausible account of the reading process.

Learning to Read

Research on reading, construed as a process of interpreting printed symbols, has a relatively long history in psychology. Scientific research on the psychology of reading began at least a century ago with the work of Cattell (1886). Other early scholars included Huey (1908/1968) and Buswell (1920). Fueled in part by its obvious relevance to a central educational task of schools, research on reading has continued in an almost unbroken line. Much of this research was stimulated by and played a role in a long-standing debate over ways to teach reading: a word recognition emphasis versus a contextual meaning emphasis; direct instruction in the grapheme–phoneme mappings of alphabetic languages (i.e., on phonics) versus focusing on words as visual wholes.

In this chapter, I consider these pedagogical debates only indirectly, concentrating instead on a body of cognitive research that sets the debates in a somewhat new light. I develop two main themes: (a) the active interplay between expectations for what will appear in a text and the visual stimuli of printed words—that is the interaction of *top-down* and *bottom-up* processes in reading; and (b) the central role of automatic—that is, very fast and nonconscious—processes of word recognition. Both the top-down/bottom-up interaction and the automaticity of processing are also important aspects of many other cognitive skills. Thus, in considering the process of learning to read, I am in fact addressing issues that are central in much research on the nature of cognitive skill.

Interaction of Top-Down and Bottom-Up Processes in Reading

The description of Just and Carpenter's READER model of text processing has already introduced the notion that there is an interaction between expectations for what will appear in a text—expectations based on the representation of the text's meaning built to date—and the actual words that appear in the text. Similar interactions between expectations and actual stimuli occur in the act of recognizing words as well as in interpreting them. To the extent that expectations for what ought to appear drive the process of word recognition, reading is considered to be a top-down process. To the extent that the printed symbols drive the word recognition process, cognitive scientists speak of reading as a bottom-up process.

In a purely bottom-up view of reading, lower level processes (i.e., detecting features of letters, combining features into letters, and combining letters into words) are assumed to occur prior to and independent of higher level processes. First words are recognized, then a syntactic processing occurs, and finally a semantic interpretation is made

based on the sentence syntax. Furthermore, these processes are controlled entirely by the printed input: Word recognition precedes comprehension of meaning. By contrast, in a purely top-down conception of reading, higher level processes, such as making inferences about meaning, are assumed to control the system, and lower level processes are called into play only as they are needed. Hypotheses about the meaning of the text are generated from prior knowledge of the topic, knowledge of the specific textual context, and a minimal syntactic parsing and sampling of visual cues. Then the printed text is used to confirm or disconfirm the hypotheses. According to an extreme top-down view, comprehension of meaning precedes recognition of words, and complete encoding of separate words may not occur at all (cf. Frederiksen, 1979).

There is ample evidence that both top-down and bottom-up processes are involved in reading. Evidence of the influence of semantic context and prior knowledge—top-down effects—includes the following kinds of phenomena: Oral reading errors, even in young readers, tend to be semantically and syntactically appropriate to the context; long hesitations or misreadings occur at points in texts where there are syntactic or semantic anomalies; people are faster at pronouncing a word in context than when the same word appears in isolation, and they are faster at pronouncing words when the preceding context is congruous with the word than when it is incongruous; word recognition is also faster when the semantic category to which the word belongs has been presented in advance (e.g., *parakeet* is recognized faster after the word *bird* than after the word *mammal*). Finally, letters can be discriminated more quickly in the context of a word than in isolation or in an arbitrary string of letters. (Resnick, 1981b). There has been less effort experimentally to establish the reality of bottom-up effects in word recognition, because it seems self-evident that people must be paying some attention to actual features of the printed stimulus as they read, else the process could not properly be called reading.

Recent research on the nature of reading has focused not on whether bottom-up or top-down processing predominates, but rather, on how the two kinds of processes interact to produce both word recognition and comprehension of a written text. Rumelhart and McClelland (1981) have developed an influential interactive model of word recognition. In this model, both features of the written words and expectations about meaning cause "activation" in the brain. The two sources of activation together, through a complex system of interaction, eventually determine what word will be "seen" by the reader.

Automaticity

The interactive model of word recognition offers an explanation of how the automatic processing of printed words might occur. It is not an

accident that psychologists have been seeking to account for complex cognition in terms of processes that do not depend entirely on conscious, planned mental activity. One must assume that much processing is automatic not only because people cannot always correctly report their processing, but also because there must be compensation for humans' limited active memory capacity.

The limited capacity of human working memory is probably the earliest fact that emerged from the beginnings of cognitive psychology. In a seminal work, Miller (1956) suggested that adults have only seven "slots" (plus or minus two) for holding information in working memory, which is where active, planned processing must occur. This notion of a limited capacity for information processing is central to all cognitive science. Psychologists are no longer certain that slots in memory is the best way to describe capacity limitations, or that there is any reality to the number 7 ± 2 as the capacity of working memory. Nevertheless, all cognitive scientists agree that there is some computing work that has to go on for thought to proceed, that the capacity for doing this is limited, and that this can create a bottleneck. That is, if too much capacity is devoted to any one component of a complex learning task, then other components will suffer.

Despite this limited processing capacity, people are able to perform complicated tasks. How? There are two major mechanisms that allow people to overcome memory capacity limitations: (a) Certain components of a task become automated so that they require very little direct attention and therefore use up little working capacity, and (b) information is "chunked" so that each slot in working memory is filled with a cluster of related knowledge. The role of automatic processing in facilitating complex performances has been investigated most heavily in the context of acquiring basic reading skills. A growing research literature that contrasts good and poor readers at various stages of development is identifying particular components of reading skill that distinguish the contrasting skill groups. A consistent finding in this research is that people who read poorly (i.e., who score poorly on standardized reading comprehension tests) also are generally slower at recognizing words. It is speed, rather than accuracy of word recognition, that seems to be important. Some individuals apparently have large recognition vocabularies and adequate word recognition skills as long as they are permitted indefinite amounts of time to process each word, but they seem to proceed so slowly that they cannot effectively understand what they are trying to read.

In the interactive theories of reading such as those just examined, timing is often crucial, for several sources of information must be integrated and thus must be present in working memory at the same time. Memory capacity is also crucial. Processes that take up too much working memory capacity or too much direct attention may drive out the other processes that are needed to provide all of the necessary information simultaneously to the system. Automation of the word rec-

ognition component of reading may be necessary both for quick and timely processing of meaning and for reducing the working memory demands that allow reading to proceed smoothly.

Establishing a correlation between automatic word recognition and comprehension skill does not of itself explain how automaticity is acquired, nor does it necessarily mean that automatic recognition causes the development of comprehension skill. To the contrary, practice in reading and comprehending texts might be the cause of improved automaticity, or automaticity and comprehension skill might both depend on some other, as yet unidentified, process. A recent longitudinal study helps to limit the possibilities. Lesgold and Resnick (1983) found that children in the first grade who have large automaticity problems are very likely to have difficulties in comprehension a year or two later. Early comprehension difficulty, however, does not predict later automaticity difficulties. This asymmetric relationship allows researchers to reject the possibility that comprehension skill causes automaticity and suggests that automaticity difficulties may indeed be helping to cause difficulties in learning to comprehend written texts.

If automaticity is a prerequisite for acquiring comprehension skill, then it should be the case that training in automaticity of word recognition would produce improved comprehension. Does it? In one study (Fleisher & Jenkins, 1978) it was found that even though speeded practice can significantly increase speed of recognizing isolated words, there is no immediate transfer to comprehension. This means that comprehension skill is not ready and waiting to be "released" by improved word recognition automaticity. However, the processes of acquiring comprehension skill may nevertheless be enhanced by increased recognition speed. If that is the case, the effects on comprehension performance would be visible only after some delay, during which time reading comprehension was practiced. Psychologists do not yet know the long-term effects of training in fast word recognition. Furthermore, training that focuses only on speed, rather than on aspects of word analysis believed to function in highly skilled reading performance, may deflect learners' attention from the very features of words that allow for automated access to meaning. A current research program (Frederiksen, Warren, & Rosebery, in press) is pursuing the hypothesis that training adolescents with very poor reading skills to quickly recognize frequently recurring spelling patterns will improve their general reading performance. These patterns are the building blocks of words and according to some theorists (e.g., Venezky & Massaro, 1979) are the units in reading that correspond directly to meaning.

Developing Mathematical Competence

When we turn to other domains of intellectual competence, many of the same themes as have been noted for natural language understanding

again emerge as central. Consider mathematics. Understanding mathematics, like understanding natural language, requires that people have a certain number of particularly powerful schemata that are used as prototypes to interpret specific expressions and situations. Further, inference processes are central in both learning mathematics and in solving mathematical problems, just as they are in reading and writing natural language texts.

In the case of mathematics, however, there is a special problem of linking symbols to their referents. Like natural language texts, mathematical expressions and mathematical procedures have both a syntax and a semantics. That is, they obey rules of "well-formedness" that are equivalent to the grammar of sentences or the rhetorical structures that constrain the more global forms of texts. In mathematics, as in formal logic, there are complex rule systems for manipulating expressions that ensure that new expressions constructed in the course of solving problems or performing algorithms will be syntactically correct. So much attention is paid to these syntactic properties of mathematics in the ordinary course of teaching and learning mathematics that people sometimes treat mathematical expressions as if they were nothing but strings of syntactically well-formed symbols.

But mathematical expressions also have a semantics—they refer to something external to themselves. These mathematical referents are quantities and relations, and it is these quantities and relations that are in fact manipulated when one performs operations with mathematical symbols. People rarely, if ever, think about natural language sentences as if they were simply sets of syntactically well-structured character strings. Instead, people treat language automatically as a way of referring to an external situation. In mathematics, by contrast, people sometimes treat mathematical expressions as if they were divorced from any referent, and this causes difficulty in learning mathematics for many people. At the same time, to be skilled in mathematical thinking requires that the person be able to manipulate the symbol system fluently. There is thus a special set of problems that arise when one analyzes mathematics as a domain of cognition and learning.

Implicit Understanding of Mathematical Principles

I begin with evidence of the role that organizing schemata have in mathematics learning. There is growing evidence that children, and uneducated adults as well, possess considerably more knowledge of certain mathematical principles than is habitually ascribed to them. This understanding is evident most typically in the kinds of informal arithmetic methods that they use. When such methods have not been taught, either formally in school or informally in the culture, they can be used to infer the kind of underlying understanding that people have of mathematical principles. Herbert Ginsburg and his colleagues (Ginsburg, 1977, 1983; Houlihan & Ginsburg, 1981) used a variety of interview

methods to document a wide range of numerical problem-solving procedures that are used by young children and by adults in unschooled cultures who have not had formal instruction in specific arithmetic routines. Similar kinds of invented procedures have also been documented by other investigators using laboratory methods of research.

Invented counting procedures. The earliest and apparently most frequent way that young children solve arithmetic problems if they have not memorized the answer is to use some form of counting. This may be "counting in the head," rather than overt counting of physical objects, as has been demonstrated in a number of studies of mental addition and subtraction. Groen and Parkman's (1972) research is the point of reference for work on simple mental calculation. They tested a family of process models for single-digit addition. All of the models assumed that a "counter in the head" could be set initially at any number, then incremented a given number of times, and finally "read out" (see Figure 2). The specific models differed in where the counter was set initially and in the number of increments-by-one required to calculate the sum. For example, the counter can be set initially at zero, the first addend counted in by increments of one, and then the second addend counted in by increments of one. If one assumes that each increment needs about the same amount of time to count, then someone doing mental calculation this way ought to show a pattern of reaction times in which time varies as a function of the sum of the two addends. This has become known as the *sum* model of mental addition.

A somewhat more efficient procedure begins by setting the counter at the first addend and then counting in the second addend by increments of one. In this case—assuming that the time for setting the counter is the same regardless of where it is set—reaction times would be a function of the size of the second addend. A still more efficient procedure starts by setting the counter at the larger of the two addends,

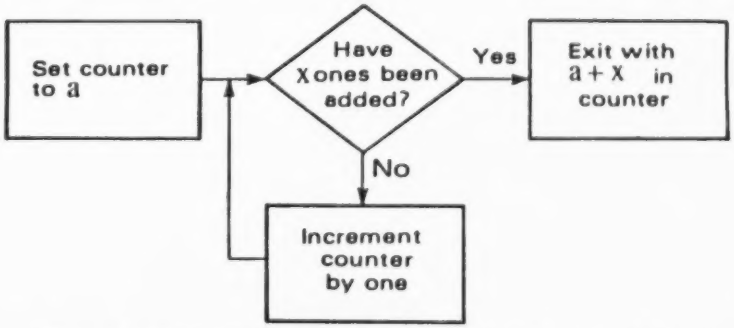

Figure 2. Schematic model for mental counting in arithmetic. Adapted from Groen and Parkman (1972) by permission.

regardless of whether it is the first or the second, and then incrementing by the smaller. Obviously, this requires fewer increments. Because this procedure produces reaction times that are a function of the size of the minimum addend, it has become known as the *min* model.

Groen and Parkman evaluated these models (along with some others that were logically possible but psychologically implausible) by comparing the predicted and observed patterns of reaction times for each model. They found that the reaction times of children as young as first-graders fit the predictions for the min procedure. Figure 3 shows a characteristic data plot. Note that problems with a minimum addend of 4 cluster together and take longer than problems with a minimum addend of 3, and so on. Subsequently, the prevalence of the min model has been confirmed in studies that have extended both the range of problems and the children studied from those aged 4½ or so to those aged 9 or 10 (Groen & Resnick, 1977; Svenson & Broquist, 1975; Svenson & Hedenborg, 1979; Svenson, Hedenborg, & Lingman, 1976).

Counting models have also been applied to other simple arithmetic tasks, especially subtraction (Svenson et al., 1976; Woods, Resnick, &

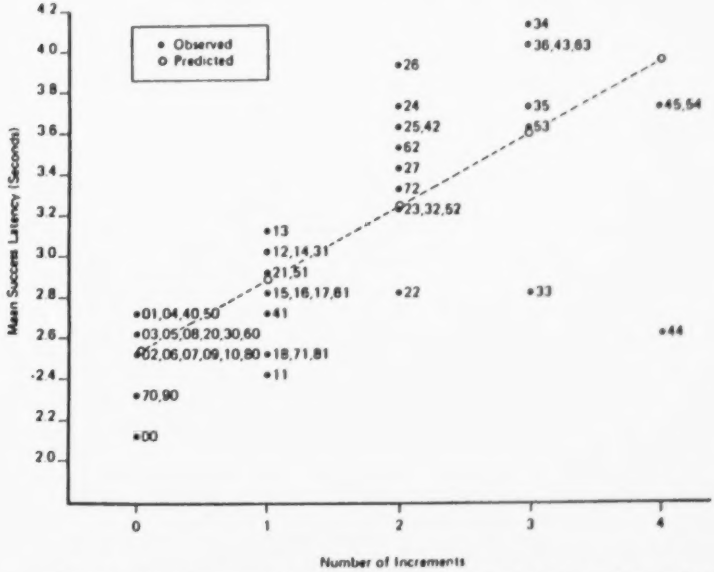

Figure 3. Reaction times for first graders solving addition problems. Pairs of numbers listed beside dots stand for single-digit addition problems. For example: 70,90 means 7 + 0, 9 + 0. Dots indicate average reaction times for adding each pair of numbers.

Groen, 1975) and addition with one of the addends unknown (Groen & Poll, 1973). In the case of subtraction, at least three mental counting procedures are mathematically correct. One procedure, *decrementing*, would involve initializing the counter in the head at the larger number (the minuend) and then decrementing by one as many times as indicated by the smaller number (the subtrahend). In the decrementing model, reaction times would be a function of the smaller number. In the second procedure, *incrementing*, the counter would be initialized at the smaller of the two numbers and be incremented until the larger number is reached. The number of increments then would be read as the answer. Reaction times for this incrementing model would be a function of the remainder, the number representing the difference between the minuend and subtrahend. In a particularly efficient procedure, *choice*, either the decrementing or the incrementing process is used for subtraction, depending upon which required fewer steps on the counter. Reaction times would be a function of the smaller of the subtrahend or the remainder. This choice model is what most primary school children use, although a few second-graders use the straight decrementing model.

It is always risky to attribute complex processes such as min and choice to people entirely on the basis of their reaction time patterns. For this reason, it is important to ask if any converging evidence exists that points to the reality of mental counting procedures. Observations of overt counting-on strategies for addition by several investigators (e.g., Carpenter, Hiebert, & Moser, 1981; Fuson, 1982; Steffe, Thompson, & Richards, 1982) suggest that the counting presumed in these models is real. Furthermore, Svenson and Broquist (1975) interviewed their subjects after each timed trial and found that on about half of the problems children reported counting-up from the larger number (by ones or in larger units).

Invented regrouping procedures. The existence of privileged, particularly well-learned number *facts* is the basis for another class of invented procedures that emphasize the regrouping of quantities. Typically, among young children and unschooled adults, not all number facts are equally well known. Those involving smaller numbers are better known than those involving larger ones, people tend to know addition facts better than subtraction facts, and more people know certain privileged types of facts: doubles, those involving the addition of 1 or 2, possibly those involving the addition or subtraction of 5, and, for older children and certain cultures, those facts involving 10 and multiples of 10. A person who cannot easily retrieve $3 + 5 = 8$ as an addition fact might regroup the problem to take advantage of a known doubles fact and would solve instead $(3 + 3) + 2$. In cultures that use a decimal notation and counting system, regrouping patterns often take special advantage of the decimal structure. Here is an example, complete with a characteristic error, from an interview study of the development of decimal number understanding:

E: Can you subtract 27 from 53?

S (an 8-year-old): 34.

E: How did you figure it out?

S: Well, 50 minus 20 is 30. Then take away 3 is 27 and plus 7 is 34.[1] (Resnick, 1983, p. 131)

Sometimes regrouping around privileged number facts and counting are combined. For example, Resnick and Omanson (in press) have used reaction-time methods to document a procedure that some children use to add a one-digit number to a two-digit number. The procedure is called *min of the units* because when people use it, their pattern of reaction times is a function of the smaller of the two units digits[2] in the numbers to be added. The person using this method decomposes the two-digit number into a tens component and a units component, then recombines the tens component with whichever of the two units digits is larger. The mental counter is set to this reconstituted number and the smaller units digit is counted in increments of one. For example, for 23 + 9, the counter would be set at 29 and then incremented 3 times to a sum of 32. The regrouping of numbers to take advantage of well-known number facts is also characteristic of people who are exceptionally good at complicated mental arithmetic (e.g., multiplying); however, such individuals have a much wider store of well-learned, privileged facts and show much more flexibility in regrouping to use the facts than do young children or unschooled adults.

Understanding implicit in invented procedures. Research of the kinds just described has now established that people use a considerable variety of invented arithmetic strategies. A concomitant step has been to show, through appropriate analyses, the kinds of understanding of mathematical principles that underlie these inventions. The first systematic effort along these lines appeared in Gelman and Gallistel's (1978) work on the nature of counting competence in very young children. They used a number of aspects of preschool children's performance to establish the fact that the children know implicitly—although they are unable to verbalize—five principles.

- The one-to-one principle: Each item in an array must be tagged with one and only one unique tag
- The stable-order principle: The tags used must be drawn from a stably ordered list

[1] The error is in the confusion about which of the right-hand digits is to be added and which is to be subtracted.

[2] In decimal notation, the digit to the far right represents "units," one multiplies it by 1 to obtain its value. The next digit represents "tens," one multiplies it by 10 to obtain its value. Subsequent digits represent "hundreds," "thousands," and so forth.

- The cardinal principle: The last tag used for a particular count represents the cardinal number of the array
- The abstraction principle: Any set of items may be collected together for a count
- The order-irrelevance principle: The order in which items in a set are tagged is irrelevant.

Greeno, Riley, and Gelman (1984) have described the counting principles as a form of *conceptual competence* that can be inferred from the *performance competence* that children exhibit on a range of counting tasks. Performance competencies are granted when a child can assemble a set of procedures that produce a performance that adheres to the conceptual principles. Conceptual competence is most clearly revealed when a new variant of a procedure must be invented. For example, when the children in Gelman's study were given an array of objects to count and told to "make this one (an object in the middle of a straight-line array) number one" or "make this one (the object in the normally first-counted position in the array) number three," the children adjusted the order in which they touched the objects but not the order in which they said the numbers, and they still touched each object only once. These children thus clearly demonstrated command of the order-irrelevance principle, the stable-order principle, and the one-to-one principle.

Data and analyses of this kind make it possible to articulate the presence of implicit knowledge and hence circumvent the need to have people state their knowledge before granting them an understanding of principles. The role of conceptual understanding that is implicit in invented procedures is also revealed in work done by Neches (1981; also Resnick & Neches, 1984). In this work Neches attempted to provide a formal account of the way that children invent the min addition procedure (described earlier) of counting on from the larger of the two addends. Neches has constructed a computer simulation program that begins by counting up both addends (essentially the sum procedure); it then modifies itself so that after a number of trials, it performs the min procedure of counting on from the larger number. To do this, the program must "discover" that setting the counter to a number will always yield the same thing as counting the objects specified by a number (a form of quantity conservation), and that it does not matter which number is set in the counter and which is added in (a form of commutativity).

Neches's program makes these discoveries by continually inspecting its own performance and applying a small set of procedure-changing heuristics. Although the final version of the program cannot be said to "know" about commutativity in the sense of explaining it, it behaves as if it understood commutativity, and it does so on the basis of its own knowledge construction without having been "told" about commutativity. Neches's program is a plausible theory—but not the

only possible one (cf. Baroody & Gannon, 1984; Resnick, 1983)—that explains how children might invent the min procedure and what understanding it is appropriate to grant them on the basis of that invention.

Schemata in mathematical problem solving. Another source of evidence demonstrating the role of implicit knowledge in children's early mathematical performances comes from research on how children solve simple arithmetic problems given in story form. Research in several countries has demonstrated great regularity in the kinds of addition and subtraction problems that are hardest to solve (Carpenter & Moser, 1982; Nesher, 1982; Vergnaud, 1982). Several analyses of this cumulative body of data have converged on an explanation of these regularities that attributes to children an understanding of the principle that mathematicians call the "additive composition of number." This principle maintains that numbers are composed of other numbers, that the number 7, for example, is not only the cardinality of the set that one can count by tagging objects up to 7, but also a composition of 1 and 6, 2 and 5, and so forth (Resnick, 1983). In the analyses of story problems, additive composition is attributed to children in the form of a *part–whole schema* (Figure 4). The schema specifies that any quantity (the whole) can be partitioned into the parts as long as the combined parts neither exceed nor fall short of the whole. By implication, the parts make up or are included in the whole. The part–whole schema thus provides an interpretation of number that is quite similar to Piaget's (1941/1965) definition of an operational number concept.

Figure 4 shows how the fundamental part–whole relation underlies several classes of story problems as well as number sentences. In each problem the whole is coded as a dot-filled bar, whether it is a given quantity or the unknown quantity. Similarly, each part is uniquely coded. The relation between parts and whole for all the problems, including the number sentences, is shown in the center display. Any bar can be omitted and thus become the unknown. Although number sentences and the given words of story problems cannot be mapped directly onto one another (Nesher & Teubal, 1975), each can be mapped directly onto a more abstract part–whole representation such as the bars shown here. The part–whole schema thus provides an interpretive structure that can permit the child either to solve certain more difficult problems directly by the methods of informal arithmetic or to convert them into number sentences that can then be solved through procedures taught in school.

Riley, Greeno, and Heller (1983) have developed a set of computational models that explain differences in the difficulty of solving certain kinds of addition and subtraction story problems. These models suggest that it is the application of the part–whole schema that makes it possible to solve difficult classes of story problems that children usually cannot solve until their second or third school year. These include set-change problems with the starting set unknown (e.g., "John

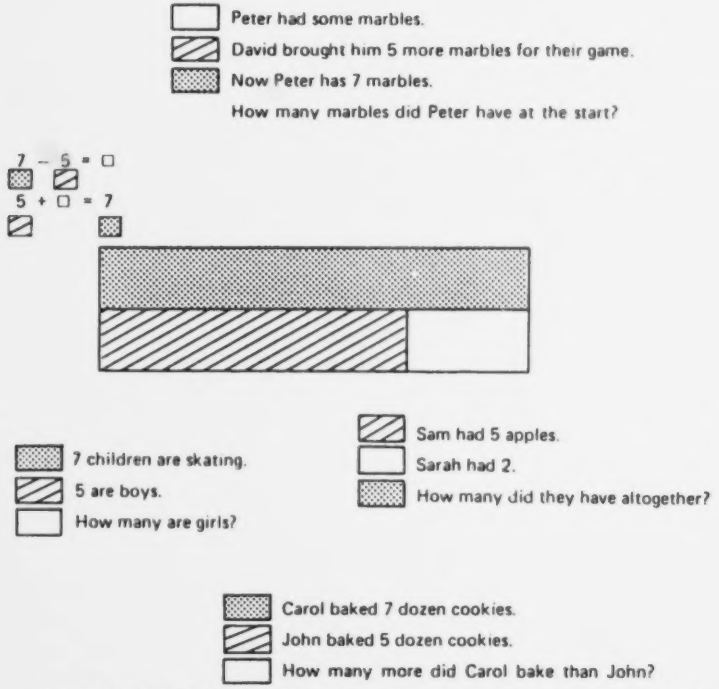
Peter had some marbles.
David brought him 5 more marbles for their game.
Now Peter has 7 marbles.
How many marbles did Peter have at the start?
7 − 5 = □
5 + □ = 7
7 children are skating.
5 are boys.
How many are girls?
Sam had 5 apples.
Sarah had 2.
How many did they have altogether?
Carol baked 7 dozen cookies.
John baked 5 dozen cookies.
How many more did Carol bake than John?

Figure 4. Mapping of stories and number sentences to a concrete model of Part–Whole. Reprinted from Resnick (1983) by permission.

had some marbles. Michael gave him 4 more. Now he has 7. How many did he have to start?") and various kinds of comparison problems (e.g., "John has 4 marbles. Michael has 7. How many more does Michael have than John?"). For these problems, solution proceeds by mapping the statements in the problem to the slots of the part–whole schema. The numbers in the problem are assigned to either *part* or *whole* status, and this permits the unknown to be clearly identified as a part or a whole.

An alternative story problem model by Briars and Larkin (1984) solves some of the more difficult problems by constructing a mental script that reflects real-world knowledge about combining and separating objects, rather than abstract part–whole relations. The script describes the actions in the story and allows the system to keep track of the sets and subsets involved. Yet in Briars and Larkin's model, too,

it is possible to solve unknown-first problems only by using a part–whole schema. Both theories show that schematic knowledge about the reference situations for mathematical problems is essential at even the simplest level of problem solving.

The Pervasiveness of Inference: Invented Errors

Side-by-side with the accumulating evidence of implicit understanding shown in informal arithmetic performances, there is equally compelling evidence of the presence of persistent and systematic errors. In fact, documenting systematic errors exhibited in the course of learning procedures is a major and pervasive feature of recent research on learning. These systematic errorful procedures are also invented by learners, but unlike those procedures described in the preceding section, these do not reflect understanding of mathematical principles. Although systematic errors in arithmetic procedures have been documented for several different parts of the school mathematics curriculum, the two that have received the most careful analysis by cognitive scientists are subtraction with borrowing and algebra. These two example domains provide contrast in detail, but they support each other with respect to the fundamental processes that seem to be involved.

In their analysis of subtraction, J. S. Brown and R. R. Burton (1978; Burton, 1982) have constructed an extensive catalog of incorrect procedures that are used by children for written subtraction with borrowing. These incorrect procedures are variants of the correct ones; they are analogous to computer algorithms with "bugs" in them and have therefore been christened "buggy algorithms." A finite number of bugs, which in various combinations make up several dozen buggy algorithms, have been identified for subtraction. Figure 5 shows a few of the most common buggy algorithms identified in this research.

These examples show that the results of buggy calculations tend to "look right": Everything is organized into columns, there is only one digit in each column, there are numbers crossed out and small digits handwritten in the conventional places, and so forth. Buggy algorithms thus look rather sensible and often contain only small departures from the correct algorithms. It appears that the buggy procedures are constructed by children when they encounter an arithmetic problem for which they have an approximate, but incomplete, rule. Rather than giving up, these children try to patch and repair the rule so that it appears to work.

J. S. Brown and K. VanLehn (1982) have developed a formal theory, in the form of a computer simulation, of the origin of bugs in arithmetic. The program invents the same bugs that children do, but not a large number of other logically possible ones. It thus constitutes a theory of the kinds of knowledge and processes that children use when they

1. **Smaller-From-Larger.** The student subtracts the smaller digit in a column from the larger digit regardless of which one is on top.

326 −117 211	542 −389 247

2. **Borrow-From-Zero.** When borrowing from a column whose top digit is 0, the student writes 9 but does not continue borrowing from the column to the left of the 0.

6 ⁹0,2 −437 265	8 ⁹0,2 −396 506

3. **Borrow-Across-Zero.** When the student needs to borrow from a column whose top digit is 0, he skips that column and borrows from the next one. (Note: This bug must be combined with either bug 5 or bug 6.)

7 ⁸0,2 −327 525	8 ⁷0,4 −456 308

4. **Stops-Borrow-At-Zero.** The student fails to decrement 0, although he adds 10 correctly to the top digit of the active column. (Note: This bug must be combined with either bug 5 or bug 6.)

703 −678 175	604 −387 307

5. **0 − N = N.** Whenever there is 0 on top, the digit on the bottom is written as the answer.

709 −352 457	6008 − 327 6321

6. **0 − N = 0.** Whenever there is 0 on top, 0 is written as the answer.

804 −462 402	3050 − 621 3030

7. **N − 0 = 0.** Whenever there is 0 on the bottom, 0 is written as the answer.

976 −302 604	8 ⁹8,6 −409 407

8. **Don't-Decrement-Zero.** When borrowing from a column in which the top digit is 0, the student rewrites the 0 as 10, but does not change the 10 to 9 when incrementing the active column.

¹⁰70,2 −368 344	¹⁰⁰0,5 − 9 1106

9. **Zero-Instead-Of-Borrow.** The student writes 0 as the answer in any column in which the bottom digit is larger than the top.

326 −117 210	542 −389 200

10. **Borrow-From-Bottom-Instead-Of-Zero.** If the top digit in the column being borrowed from is 0, the student borrows from the bottom digit instead. (Note: This bug must be combined with either bug 5 or bug 6.)

702 −248 454	508 −489 109

Figure 5. Descriptions and examples of Brown and Burton's (1978) common subtraction bugs. Adapted from Resnick (1982) by permission.

invent buggy subtraction algorithms. The repair theory program is a "generate and test" problem-solving routine of the kind that characterizes many successful performances in other domains (cf. Simon, 1976).

According to the theory, buggy algorithms arise when the child encounters an arithmetic problem for which his or her current algorithms are incomplete or inappropriate. The child, trying to respond, eventually reaches an impasse, a situation for which no action is available. At this point, the child generates a candidate *repair* by calling on a list of actions to try when a standard action cannot be used. The repair list includes strategies such as performing the action in a different column, skipping the action, swapping top and bottom numbers in a column, and substituting an operation (such as incrementing for decrementing).

The outcome generated through this repair process is then checked by a set of *critics*. The critics inspect the resulting solution for conformity to some basic criteria such as no empty columns, only one digit per column in the answer, only one decrement per column, and the like. Note that the generate-and-test problem solution calls on no knowledge about the quantities that the numerical symbols represent. This is a crucial characteristic of buggy arithmetic, and one that I will return to.

Several researchers (Carry, Lewis, & Bernard, 1980; Davis, 1983; Greeno, 1983; Sleeman, 1982) have studied the errors (often called *malrules*) that students make in algebra. When students apply the rules of transformation that are the basic tools of algebraic problem solving, these investigators have shown (a) that many errors are made by beginners as a result of either incorrect rules or incorrect applications of correct rules; (b) that these errors persist for a long time, showing up occasionally even among expert algebra performers; (c) that there is great systematicity in which errors appear in different students (i.e., only a small number of the logically possible algebra errors actually tend to be made); and (d) that there is, nevertheless, a lack of stability in the performance of any given individual (i.e., learners do not always apply the same algebra malrule even in what is, to the experimenter at any rate, the same situation).

The best developed theory to date that explains how these malrules are invented is one by Matz (1982). Matz's theory, like the Brown and VanLehn theory of subtraction bugs, is expressed as a simulation program that invents the malrules that were observed in algebra solutions; other possible malrules are not invented. Matz proposes that children learning algebra construct prototype rules from which they extrapolate new rules. Although the results are malrules, both the construction of the prototypes and the extrapolation follow regular principles. An example appears in Figure 6.

The initial rule is the distribution law that is typically taught in a beginning algebra course. From this correct and specific rule, a prototype is created by generalizing over the operator signs. That is, the prototype specifies not that multiplication ($\times$) can be distributed over addition (+), but that any operator ($\square$) can be distributed over any other operator ($\triangle$). From this prototype, new but incorrect distribution

1. The correct rule as taught:

$$a \times (b + c) = (a \times b) + (a \times c)$$

2. Prototype created by generalizing over operator signs:

$$a \square (b \triangle c) = (a \square b) \triangle (a \square c)$$

3. Incorrect rules created from the prototype.

$$a + (b \times c) = (a + b) \times (a + c)$$

$$\sqrt{b + c} = \sqrt{b} + \sqrt{c}$$

Figure 6. Example of the formation of an algebra malrule. Based on Matz's (1982) theory.

rules can be constructed by substituting specific operations for the generalized operators in the prototype. The elegance of Matz's model does not prove that it is a correct theory of the origin of algebra malrules in human learners, but it does establish the conditions for an ongoing discussion of the nature of malrule invention (see, e.g., Sleeman, 1982) that is specific about the knowledge and processes likely to be involved. Such a discussion is therefore useful both for understanding difficulties in mathematics learning and for explicating general principles of cognitive acquisition.

The prevalence of buggy algorithms and malrules in mathematical learning points to a pervasive feature of human cognitive functioning. It is natural to seek meaning and to draw inferences. People will do this on the basis of whatever knowledge they have available—even if it is incomplete or incorrect. For this reason, perfectly good inferential and reasoning processes will sometimes produce errors.

Linking Symbols and Their Referents

Buggy algorithms and algebra malrules also point to a special difficulty that must be overcome whenever formal representational and rule systems are part of the subject matter to be learned. This relation between formal systems and intuitive or informal ones is most evident in math-

ematics, where specialized notational systems and rules for manipulating them amount to a new language, complete with grammatical rules, that must be mastered. The difficulty is that for a language to function appropriately, its grammar and formal rules must work in concert with its referential system. In technical terms, the syntax and the semantics of a language system must function together so that sentences are "well-formed" (i.e., they obey the grammar of the language) and at the same time the referents of the sentences are clear (i.e., they maintain the semantics of the language). In natural languages, this coordination of syntax and semantics seems to occur without any special work or attention on a human learner's part. In an earlier section of this chapter, I discussed models of text comprehension that show that people reading or listening to a natural language text naturally and quite automatically build up a representation of what the text refers to. In mathematics, syntax and semantics sometimes become separated.

A reconsideration of buggy subtraction can make this point more clearly. I have already noted that subtraction bugs seem not only to respect the syntax of written arithmetic, but also to disobey constraints that would be apparent if the quantity referents were being kept in mind. For example, consider the second bug, borrow-from-zero, in Figure 5. At a strictly symbolic level, this procedure seems a reasonable response to encountering a zero in the course of borrowing. The zero is changed to 9, which is a familiar result of borrowing when zeros are present. However, the bug violates the fundamental principle that the total quantity in the minuend must be conserved during a borrow. Interpreted semantically—that is, in terms of quantities rather than simply manipulations of symbols—a total of 100 has been added to the minuend, 10 in the units column, and 90 in the tens column, with no compensating decrement in the hundreds column. The next bug, borrow-across-zero, shows a similar disregard for the need to conserve the minuend quantity. The bug respects the syntactic rules for symbol manipulation that require that a small "1" be written in the active column and that some other (nonzero) column be decremented. The bug violates the conservation principle, however, by removing 100 from the hundreds column but returning only 10 to the units column.

This informal analysis is supported by reexamining Brown's and VanLehn's repair theory of the origin of subtraction bugs. The repair theory program produces bugs by generating repairs and checking them against critics. All of the critics in the program are syntactic in nature; that is, they reflect rules for symbol manipulation, but they do not embody any knowledge of principles of quantity. The fact that repair theory matches human performance by inventing only the bugs that children do and not other logically possible bugs suggests that children represent subtraction to themselves as sets of rules for transforming symbols without reference to the quantities that these symbols in fact are meant to represent.

Further evidence that symbols take on a life of their own, apart from the quantities that they represent, comes from a training study (Resnick & Omanson, in press) in which children who had been diagnosed as using buggy algorithms were taught the correct principles of subtraction just described (e.g., conserving the minuend quantity). This teaching was done in a form that insured that the children's knowledge of the principles was in fact linked by them to the steps in the algorithm they were being taught. Detailed interviews established with considerable certainty that some of the children had fully understood the principles and their application to written subtraction algorithms. Nevertheless, as soon as they returned to a situation of routine calculation performance, half of those children who understood and could apply the principles returned to their buggy algorithms. That is, they did subtraction in a way that violated principles they clearly knew. This is further evidence of a tendency within mathematics for the syntactic system to become separated from its semantic referents even when the necessary knowledge of principles is in fact present in the individual. Similar evidence of this tendency can be found in symbolic logic and algebra problem solving. When this major problem in human cognitive functioning is better understood, cognitive scientists may be able to eliminate it through changed forms of instruction and teaching.

Problem Solving, Intelligence, and Learning Abilities

Problem solving is in a very real sense the birthplace of cognitive science. Efforts that began in the late 1950s and culminated in 1972 with the publication of *Human Problem Solving* by Newell and Simon showed how intelligent computer programs could reason and solve problems, not by doing the kind of dumb, exhaustive searching of a very large memory that was assumed to be the principal capability of computers, but by using strategies to analyze a problem situation and to select actions most likely to advance toward a specified goal. What is more, evidence was developed in the course of these efforts to show that the behavior of the programs using these methods matched in significant ways the behavior of humans working on similar problems. That is, when humans solved the problems, "thinking aloud" as they worked, they showed particular points of hesitation, backtracking, and insight, and they made typical kinds of errors. The computer programs often showed the same kinds of hesitation, backtracking, and errors. The processes built into the programs could therefore be supposed to be functioning in humans as well, although they could not be directly observed in humans.

Of course, programs never matched human performance exactly, and investigators were careful to specify what parts of human perfor-

mance were not well explained by the computer-expressed theories. Clear identification of these nonmatches was a powerful spur to successive stages of research and theory development. I have already provided an example of how a nonmatch between a theory and human performance reveals the need for a different level of theory in the Kakra and Gum story. The microcoherence-building model of Kintsch and VanDijk would have accepted Kakra's marriage after his death without hesitation, but human readers (who are not thinking about ghost marriages) immediately reject it as impossible. This kind of mismatch made it clear that the microcoherence model alone could not account for how people understood texts, and suggested that a macrocoherence model was needed as well. Noting mismatches and using them to direct further research is characteristic of all of the work that uses computer simulation as a form of theorizing about human thinking.

General Heuristics in Problem Solving

The early research on problem-solving focused on a set of puzzle-like tasks well-suited to initial efforts. The tasks studied included theorem proving in symbolic logic (a task in which all legal expressions and all allowable transformations are specified and the problem-solver must show how it is possible to derive a target expression from a given expression), cryptarithmetic (a decoding puzzle in which letters of the alphabet stand for digits, and a solved arithmetic problem sets constraints on which letters can have which digit values), a variety of artificial problems (such as the Tower of Hanoi, or Missionaries and Cannibals), and, finally, chess. With the exception of chess, which has been shown to depend heavily on extensive knowledge of chess positions, chess moves, and their likely effects, all of the problems studied depended only minimally on knowledge beyond what could be supplied in the experimental situation itself. In these knowledge-poor task environments, cognitive scientists focused their efforts on identifying general processes of problem solving.

Several strategies of problem solving that could be properly called general methods were identified and elaborated in the course of this work. I have already mentioned some of them: For example, the generate-and-test method is usable whenever there is a limited set of possible operators or objects that can be tested to see whether they meet a current goal. Another general method, recurring in many problem-solving models, is means–ends analysis, a kind of general heuristic that reduces the length of search through long-term memory. In means–ends analysis, the problem solver compares the current situation with the goal situation and identifies specific differences between them. A subgoal is then set to reduce a difference that has been identified. (A special set of heuristics governs which subgoal to work on first.) Then

a search is conducted to find the operator that will reduce the identified difference. This is a very abbreviated search, because operators are assumed to be organized in memory according to the goals they can serve. It is important to note that means–ends analysis assumes a system that has intentions (goals) and acts on them: It is capable of analyzing its situation and planning its actions on the basis of goals, albeit in a very restricted domain. The General Problem Solver (GPS) was one of the first programs to instantiate all of these general methods in a system that solved symbolic logic problems (see Ernst & Newell, 1969; Newell & Simon, 1972).

More recent work in problem solving has, along with the rest of cognitive psychology, become much more focused on performance in information-rich domains. Many of the basic strategies of heuristic search, subgoal formation, and the like turned out to be relevant for these domains as well. But it has also proved necessary to attribute to the problem-solver, whether human or artificial, specific and organized knowledge about the domain in which problem solving is to take place. Some of the best demonstrations of the role of organized knowledge in problem solving have come in recent research that compares novices and experts in physics as they solve the kinds of problems that are characteristically given as exercises in college-level physics textbooks. In these studies, good beginning students have been compared with advanced students or teachers. The studies show that one's initial understanding, even of a simple textbook problem, depends upon one's level of knowledge in the field.

In one study (Chi, Feltovich, & Glaser, 1981), novices and experts were asked to sort physics textbook problems on any basis they wished. Novices grouped problems on the basis of the kind of apparatus involved (lever, inclined plane, balance beam, etc.), the words used in the problem statement, or the visual features of the diagram presented with the problem. Experts classified the same problems on the basis of the underlying physics principle that was needed to solve the problem (e.g., energy laws, Newton's Second Law). Some typical novice classifications are shown in Figure 7; the contrasting expert classifications are shown in Figure 8. Clearly, novices are affected more by the way the problem is presented, whereas experts bring their own knowledge of important principles to bear in a way that reshapes the problem, usually into a more solvable form. This is much like the way in which good readers use their past knowledge about the topic or the form of discourse to impose a useful structure on a text, while beginning readers are much more victimized by poorly written material or indirect forms of expression.

Initial differences in the ways that experts and novices sort and classify problems are only the beginning, however; the process of solution is also different. What novices usually do is to translate the given information directly into formulas. They then work on the formulas

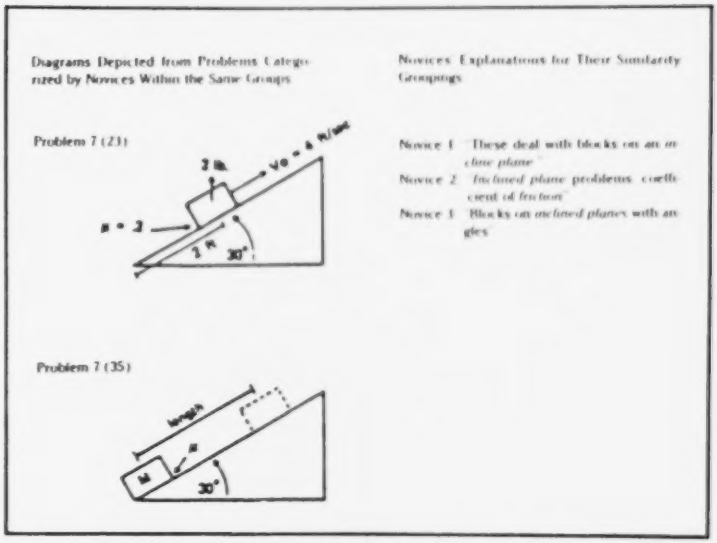
Diagrams Depicted from Problems Categorized by Novices Within the Same Groups
Novices' Explanations for Their Similarity Groupings
Problem 7 (23)
2 lb
Vo = 5 ft/sec
μ = .2
2 ft
30°
Novice 1: "These deal with blocks on an incline plane."
Novice 2: "Inclined plane problems, coefficient of friction."
Novice 3: "Blocks on inclined planes with angles."
Problem 7 (35)
length
M
μ = 0
30°

Figure 7. Diagrams of physics problems categorized by novices as similar, and samples of three novices' explanations for their similarity. Reprinted from Chi, Feltovich, and Glaser (1981) by permission.

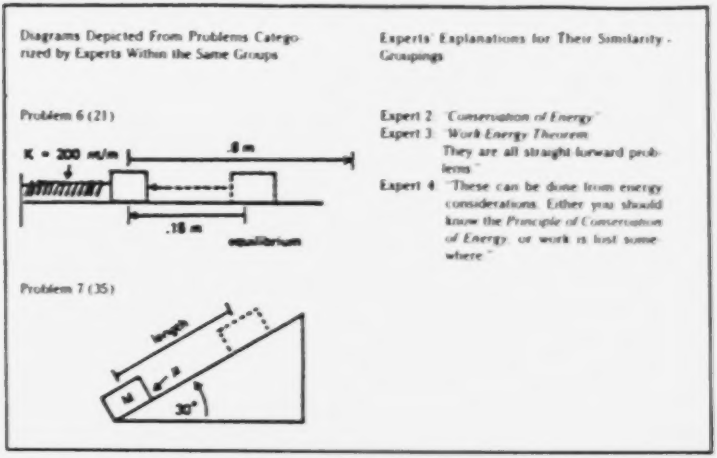
Diagrams Depicted From Problems Categorized by Experts Within the Same Groups
Experts' Explanations for Their Similarity Groupings
Problem 6 (21)
K = 200 n/m
.8 m
.15 m
equilibrium
Expert 2: "Conservation of Energy."
Expert 3: "Work Energy Theorem. They are all straightforward problems."
Expert 4: "These can be done from energy considerations. Either you should know the Principle of Conservation of Energy, or work is lost somewhere."
Problem 7 (35)
length
M
μ = 0
30°

Figure 8. Diagrams of physics problems categorized by experts as similar, and samples of three experts' explanations for their similarity. Reprinted from Chi, Feltovich, and Glaser (1981) by permission.

using rules of algebra. Experts, by contrast, do not begin by translating into formulas. Instead, they work for a while on reinterpreting the problem and specifying the various objects and relations in the situation described. They may draw diagrams to express these relations. By the time they are ready to write equations, the experts have virtually solved the problem. They do much less calculation than novices, at least on the simple problems studied so far in this research. Experts, in other words, construct a new version of the problem for themselves, one that accords with the information actually given, but one that is reformulated in terms of general principles and laws that make the solutions more apparent.

I have used physics research to illustrate the kinds of differences that have been observed in the problem solving and reasoning of novices and experts. But these differences occur in other domains as well. Similar differences have been found in tasks as divergent as interpretation of x-ray photographs by physicians, arithmetic problem solving by elementary school children, and economic planning by political scientists. In each case the more expert problem solver does not simply respond to the problem in the terms presented, but instead reinterprets it in ways that reveal an underlying structure that makes the solution sometimes appear self-evident. The similarity of this reformulation process to the processes involved in reading comprehension described earlier is not an accident. It is a fundamental reflection of the nature of human reasoning and of the constructive character of learning and thinking.

Bottom-Up Processes in Reasoning and Problem Solving

There are other ways, as well, in which research on problem solving echoes themes that I have already discussed in this essay. For example, both top-down and bottom-up processes play a role in problem solving as they do in reading. The various heuristic strategies that I have considered up to now, such as means–ends analysis and subgoal formation, are essentially top-down kinds of processes. A system using them imposes a general plan, developed in the course of prior problem-solving experience, on the specific stimuli of the problem at hand, and the stimuli presented are interpreted in terms of this plan. Similarly, the expert physics problem solvers use their prior knowledge to reorganize the problem that is given. They are more top-down solvers than are the novices. The focus on general problem-solving methods and on the directing role of prior knowledge has led research attention away from the role of the stimuli and the problem setting itself in the solution process.

Research carried out some years ago in my own laboratory (Magone, 1977) helps to clarify the role of bottom-up or stimulus-driven

processing in problem solving. In one study, subjects were asked to solve one of the classic Gestalt problems. Two strings were suspended some distance apart in a room, and the subject was asked to tie them together—a simple assignment, except that the strings were deliberately made so short that it was not possible to reach both of them at the same time. Three classes of solutions to this problem were possible: extension (tying another long object to the string or extending the arm by using some rigid, long object to hook the string in); anchoring (holding one string down in the middle, while walking over to reach the other one); and pendulum (putting a weight on one string and swinging it toward the other). There were six objects available that could be used to help solve the problem. They were chosen so that each one tended to suggest a particular class of solution, but some could be used for more than one class. All subjects had to try to solve the problem using each of the six objects in succession.

There was an important difference in the way in which the task instructions were given that tended to evoke either top-down or bottom-up patterns of solution. Some of the subjects in the study were asked by the experimenter to use the first object to solve the problem and to use a specific class of solution. These subjects tended to use the same class of solution over and over on successive objects, sometimes even using the objects in quite unusual and difficult ways in order to stay with their preferred solution type. Another group of subjects were simply told to use all of the objects. These subjects typically picked up the objects in the order in which they fell to hand and used each object in its most characteristic way. They were, in other words, object-driven, and thus bottom-up solvers. It is of interest to note that the bottom-up solvers succeeded in solving the problem in essentially 100 percent of their tries: That is, they found a way to use every object to tie the two strings together. The top-down solvers, by contrast, had a somewhat lower rate of success: They sometimes failed to see an obvious way of using an object because they were intent on trying to make it fit into their top-down solution strategy.

Intelligence and Learning Abilities

Over the decades, one of the most provocative and difficult questions faced by psychologists is the nature of intelligence and the extent to which it can be modified through teaching or other environmental interventions. The previous descriptions of research on reasoning and problem solving lead quite naturally to the question of whether general reasoning skills can be taught and whether this might improve people's general ability to learn. These questions are by no means new ones. In one form or another, they have motivated major branches of psychological research and divided both psychologists and the public at large

(see Curtis & Glaser, 1981, for a useful discussion of the history and present status of research on intelligence).

The question that has produced the most dissent, because it is so tied to social policy, is whether differences in intelligence are inherited or are acquired as a result of differences in experience. Posed in this way, the question has no single answer, and it often evokes social and political rather than scientific responses. The understanding of intelligence becomes much more tractable when one raises more pointed questions about what constitutes intelligence and how schools and other educational institutions can best respond to individual differences. These are questions to which current cognitive research on the nature of reasoning, problem solving, and learning can successfully respond.

In so doing, there are two fundamental ways of thinking about intelligence. One is to treat intelligence as something that children (or adult learners) bring with them to an educational experience. In this view, intelligence and aptitudes set the range and limits of the learning that can be expected. This view of intelligence can be held regardless of whether one believes that intelligence is set by heredity or formed by experience. Whatever the origins of individual differences, by the time people present themselves for a particular lesson or course of study, particular aptitudes may determine what kinds of learning activities will be most successful, and established capabilities for learning will set boundaries on what can be expected in a given period of time. A large amount of recent research has centered on trying to identify the cognitive processes that are involved in various kinds of intelligent performances. This knowledge, it is hoped, will provide the basis for more effectively adapting instruction and teaching to the intellectual capacities and propensities that people bring to school with them.

The second view of intelligence and aptitude is that they are processes of thinking and reasoning that can be formed by instruction. This view need not imply a total rejection of a hereditary basis for individual differences. It only requires a belief that environmental factors, including instruction, can make some significant difference in subsequent abilities to learn and reason. Although the hope of improving intelligence and learning skills through deliberate cultivation of certain ways of thinking is an old one—witness a long history of programs for improving memory, problem solving, and other learning abilities—and one that has often been disappointed, recent research on the nature of skilled learning and performance in many domains is building a more scientific basis for these efforts. Although a magic potion for curing failures in human learning and intelligence is hardly in the offing, current lines of investigation are refocusing the issue in profitable ways, and there is room for reasonable optimism about eventual developments that will extend the limits of human learning capacities.

Intelligence Brought to School: Aptitude-Treatment Interactions

If intelligence and learning ability are viewed as capacities already formed in people before they enroll in an educational program, what education can best do is adapt to intelligence—that is, provide forms of instruction that are optimally matched to the aptitudes the individual already has. This is a very old and honorable ambition of educators, although viewpoints on effective and appropriate ways of adapting to individual differences have shifted over time in response to both political and social pressures on the educational system and to available psychological theory about the nature of individual differences.

The first deliberate effort to adapt educational offerings to individual differences was made by separating (i.e., tracking) children into groups according to their different ability levels. The general idea supporting such practices was that people differed in the speed at which they could learn and perhaps also in the highest levels of abstraction to which they could aspire. By grouping faster learners with faster and slower with slower, both groups could proceed at a pace suited to their natural abilities. This, it was proposed, would produce optimal—but not necessarily equal—outcomes for each group. This theory of grouping and tracking was totally consonant with theories of intelligence and aptitude that were dominant from the end of the 19th century through at least the 1920s. Intelligence was viewed as a largely fixed trait, hardly modifiable by experience, and as a unitary trait: *general* intelligence was what determined speed and ease of learning in all domains. Binet's intelligence test and its various offspring, some still in use today, are based on this view of intelligence, as are many of the landmark research studies on intelligence and school learning of the early part of this century (see Carroll, 1982, for this history).

Beginning in the 1930s a more differentiated view of intelligence and mental abilities became predominant among psychologists, who had, through factor analysis and related techniques, identified a variety of *differential aptitudes* in which people might vary. The new viewpoint was that specific aptitudes, rather than general intelligence, were what suited people for specific forms of learning and job performance. This view of intelligence fit well with a new social and political mood that became dominant after World War II. People began to question the suitability of an educational system that more or less permanently classified children as either fast or slow learners and thus limited the potential aspirations of those characterized as slower. Further, increasing sensitivity to ethnic and cultural variations in the American population began to produce the view that recognizing different aptitudes and approaches to learning, rather than emphasizing deficits in general intelligence, would be a more suitable way to optimize educational out-

comes. Glaser (1977) called this a shift from a selective theory of education to an *adaptive* one.

The shift to an adaptive theory of education and a more differentiated concept of aptitude and intelligence has produced a search for qualitatively different instructional treatments that would be optimally matched to learners with different characteristics. The result has been research on *aptitude-treatment interactions* (ATIs). This research seeks situations in which a given instructional treatment produces different outcomes in people of different aptitudes. Optimally, one would hope for interactions that allow one to choose a treatment for each individual that will produce the highest level of performance possible in a domain, thus eliminating overall differences in performance. In fact, such ideal interactions are almost never found.

The most typical finding in ATI research is one in which a single aptitude measure—some form of a general intelligence measure—interacts with two broad classes of instructional approaches. Highly structured treatments (e.g., careful sequencing of instructional materials, required responding at specified points, teacher control, and instructions to process in a particular way) reduce the correlation between general intelligence and achievement, whereas unstructured treatments (e.g., much student control of sequence and pace, "discovery-learning" conditions, and open-ended problem setting) maintain a correlation that favors high general intelligence students. That is, low intelligence students do better under structured conditions, which are interpreted as reducing the burden of information processing for the learner (Snow, 1976). Many theories (e.g., Cronbach, 1970) suggest that high general intelligence students should do less well under these circumstances. Only a few studies, however, show such a suppression. This may be due to the fact that the tests used to assess learning often do not permit high intelligence students to demonstrate the additional knowledge or skill that they have in fact acquired in the less-structured teaching conditions.

It should not be surprising, on reflection, that the ATI enterprise as traditionally conducted has not resulted in the kind of strong basis for adapting instruction to aptitudes that had been sought. The near-total dependence in ATI research, until very recently, on standardized tests as measures of intelligence and aptitude has meant that the research attempted to match aptitudes whose characteristics were ill-understood to instructional treatments defined only in very global terms (structured versus unstructured, for example). To break this logjam, and to discover whether there are in fact ways of adapting instruction to specific rather than general capacities for learning, it is necessary to understand better what mental processes are actually involved in the various traits called aptitudes, and what kinds of processes are actually called upon in the various instructional treatments. In other words,

cognitive analyses of both aptitudes and instructional treatments are required.

Cognitive Analyses of Aptitudes

Responding to this need, there has been considerable effort in recent years to reanalyze the constructs of intelligence and aptitude in terms of cognitive processes and constructs (Friedman, Das, & O'Connor, 1980; Resnick, 1976; Snow, Federico, & Montague, 1980; Sternberg & Detterman, 1979). Most of this new work began with traditional aptitude tests (for which there is a considerable validation history, based largely in factor-analytic research) and sought to redescribe these aptitudes in terms of current cognitive constructs and parameters. Pellegrino and Glaser (1979) have made a useful distinction between a *cognitive correlates* approach and a *cognitive components* approach to the study of intelligence. The correlates approach uses an aptitude test as a criterion measure and seeks more elementary cognitive processes that are highly correlated with the test criterion. The cognitive components approach uses the test items as tasks to be analyzed in a search for the component processes of test performance itself.

Cognitive correlates of aptitude. Much research is being done to identify basic cognitive processes that distinguish between high and low scorers on a particular aptitude test. The primitive processing parameters for study are drawn from the mainstream of basic research on cognitive processes, especially memory processes. This line of research was initiated by Hunt (1978), who suggested that verbal performance requires both the specific verbal knowledge that is called upon by the task and the exercise of certain mechanistic processes by which information is manipulated. According to Hunt's theory, individuals with less efficient mechanistic processes have to work harder at learning tasks involving verbal information. Over time this handicap produces relatively large individual differences in verbal skill and knowledge.

The theoretical argument is buttressed by data from studies that have investigated the relations between performance on laboratory information-processing tasks and scores on global measures of aptitude, such as IQ tests and college admissions tests. Although early efforts (e.g., Hunt, Frost, & Lunneborg, 1973) were attempts to find associations with quantitative as well as verbal ability, the main findings have shown correlations with tests of verbal aptitude or general intelligence measures that are heavily verbal in character.

The most robust finding in this literature reveals differences in the amount of time that various people need to access name codes in long-term memory. *Code access* time is inferred from the difference between the time it takes a person to decide whether two stimuli that look

different have the same name and the time it takes the person to decide that two other stimuli are physically identical. For example, a subject might be shown a lowercase *a* and an uppercase *A* (different physical form, but same name). On a subsequent trial, the subject might be shown two uppercase *A*s (same physical form). The subject will take longer to decide that the pair of stimuli with different physical forms nevertheless has the same name than to decide that the other pair of stimuli has the same form. The longer time is needed to access the name code in memory. The extent of the difference in the time required for the two decisions correlates with verbal aptitude. Across a number of studies, the time difference tends to increase as one moves from highly verbal university students to young adults not in a university, to normal elementary school children, and finally to mildly retarded school children (cf., Bisanz, Danner, & Resnick, 1979; Hunt, 1978). Several other tasks, all requiring speed in particular kinds of microprocesses, have also been shown to discriminate high and low scorers on verbal aptitude tests.

All told, there seems to be enough evidence of individual and age differences in primitive parameters of mental processing to make plausible Hunt's notion that small differences in mechanistic processes could cumulate over time to produce considerable differences in verbal skill and knowledge. It is important to note, however, that a large portion of the findings clearly associating these parameters with individual differences comes from Hunt's own laboratory. Wider replication is needed before strong conclusions about specific associations are drawn. A recent summary and useful critique of this research appears in Cooper and Regan (1982).

Cognitive components of aptitude. Carroll (1976) and Simon (1976) first suggested the analysis of test items as cognitive tasks, and several research programs subsequently focused on uncovering the processes that are required in actually performing the items in intelligence and aptitude tests. Perhaps the most ambitious program in terms of the range of tasks studied is Sternberg's work on what he calls a "componential analysis" of intelligence (1977a, 1977b, 1980). Sternberg's analyses begin with a specification of the components that are hypothesized to be involved in the performance of a test item. Several models are then defined that differ in the components called on, the sequence of the components, the number of times each component needs to be executed, and the manner of execution (e.g., exhaustive or self-terminating searches). These models permit predictions of reaction time and error patterns under varying conditions of stimulus structure and task presentation.

Empirical tests of models generated for analogies, for example, have identified a "best fit" model and provided estimates of which processes absorbed most of the processing time. For verbal analogies, encoding of the stimulus terms accounted for about half of the solution

time, while 30 percent of the time was spent on attribute comparison operations. For geometric analogies, attribute comparisons took longer, both as a percentage of total time (57 percent) and in absolute terms. Sternberg has extended the analysis of analogies to children, making it possible to chart developmental changes in the various components. The most important developmental observation has been that children have a greater tendency to rely on associations between the words in the analogy than to analyze all of the relations.

Other research on analogies performance is largely in agreement with Sternberg's findings on the importance of encoding. Some of the studies have analyzed the encoding process itself further, with particular attention paid to which aspects of the stimuli are encoded. For example, Mulholland, Pellegrino, and Glaser (1980) showed that in geometric analogies, individuals analyze stimuli in a systematic serial manner, so that latency of responding is a function of both the number of elements that must be encoded and the number of transformations that must be performed on each element. They found a sharp increase in both reaction time and errors when multiple transformations on multiple stimuli had to be processed, suggesting that working memory limitations are important in analogy processing. For verbal analogies, studies by Pellegrino and Glaser (1980) and Sternberg (1977a, 1977b) all show that individuals with high aptitude test scores specify more precisely the set of semantic features that relate the word pairs in an analogy, and that the extra time they spend on this process allows them to spend less time on subsequent decision and response processes.

Other test-like tasks that have been subjected to similar analysis include series completion, syllogistic reasoning and transitive inference, spatial abilities tasks such as mental rotation and visual comparison, block designs, and tasks from the Ravens Progressive Matrices test. Not all of this work has been explicitly oriented toward detecting individual differences. Instead, much has been inspired by the Piaget-generated debates over how and when various logical abilities develop in children, and over whether language or spatial representations are central (see Resnick, 1981a for a review).

It seems likely that as efforts to understand performance on such tasks proceed, individual differences will have to be considered if the data are to be sensibly interpreted. An interesting case in point is Cooper's (1980) research on visual comparison, in which subjects separate naturally into two quite different subgroups, one using a holistic and one an analytic comparison strategy. The two strategies produced very different patterns of latencies, and the groups responded in predictably different ways to variations of task instructions and of stimuli. Cooper and Regan (1982) have suggested that differences in preferred strategy for various tasks, verbal as well as visual, may account for aptitudes even more strongly than across-the-board differences in speed of basic processes. Their discussion suggests ways in which correlational and

componential approaches to the analysis of aptitude may have to be joined before really adequate theories of the nature of individual differences in test performance can be developed.

Intelligence Shaped by School: Teaching Learning Skills

If one views intelligence and aptitude as a set of capacities that are formed partly by instruction, one is led to pose two questions. First, what skills of learning are sufficiently pervasive and general (that is, not limited to specific subject matters or specific situations of application) that they warrant concerted attention as the goals of educational programs? Second, how are these skills acquired, and correlatively, how might they be most directly taught? The search for general skills of learning has been a long one, and it has been pursued from many points of view. Before proceeding to a consideration of particular skills of learning and their acquisition, it is worth pausing to ask whether it is likely that such generalizable abilities exist at all.

Skepticism about the existence of general abilities. There are two bodies of evidence, one old and established, one quite recent, that must lead to skepticism toward the claim that cognitive abilities are really very general. The first set of evidence is the repeated failure, over decades of trying, to produce convincing demonstrations of widespread transfer of learning from one domain to another. The second is evidence of the central role of specific knowledge in intelligence performance and in learning.

There has been a recurrent view that certain school subjects would "discipline the mind" and should therefore be taught not so much for their inherent value as for their value in facilitating other learning. Latin was defended for many years in these terms; mathematics and formal logic are often so defended today. Most recently, learning to program computers has been offered as a way to develop general problem-solving and reasoning abilities, appropriate even when no computers are available or applicable to the situation at hand (Papert, 1980), and a variety of courses and programs claiming to teach reasoning and problem-solving abilities have been developed and promoted (see Segal, Chipman, & Glaser, in press; Nickerson, Salter, Shepard, & Herrnstein, 1984). This view of transfer from a particularly powerful or nodal knowledge has never been supported empirically. In the 1920s, Thorndike (1922) studied transfer among school subject matters and found that it was always more efficient to study the subject of interest directly (English vocabulary, for example) than to study some other subject (Latin, for example) that "prepared" one's mind. Subsequent reviews of research on transfer of school subject matter have reconfirmed Thorndike's finding, and there is as yet no empirical evidence of transfer to other activities from specific kinds of problem-solving courses or from learning to program computers.

The second source of evidence that weakens claims for generalized abilities is the research yielding repeated demonstrations that specific knowledge plays a central role in reasoning, thinking, and learning of all kinds. Several examples of the role of specific knowledge have been developed in the course of this essay. For example, specific knowledge about the topic of a text affects the processes of language comprehension, and specific, acquired schemata underlie problem-solving performances as varied as those in primary school arithmetic and college-level physics and political science. Glaser (1984) further described evidence of the role of domain-specific knowledge in a variety of tasks that have traditionally been viewed as indicators of aptitude or intelligence.

Belief in the reality of general skills. Despite the evidence that opposes transfer and that favors the importance of domain-specific knowledge, there are some equally compelling factors that have sustained psychologists' belief in the reality of general competencies in learning. First, there is a positive correlation between almost any two cognitive performances that have ever been measured, except when tests have been specifically designed not to correlate with IQ (as, for example, certain creativity tests). This *positive manifold* is the basis for the factor analytic tradition in intelligence research: Factor analysis uses patterns of covariance to infer what various tests may have in common, and thus what the basic dimensions of human aptitude are. Tests that are positively correlated—that is, that share variance—also presumably share underlying processes. The fact that most tests correlate positively with each other, and that a general factor can always be found if the statistical methods used do not insist on completely uncorrelated factors, suggests that all tests have some processes in common. These common processes are, presumably, general abilities.

Second, when cognitive scientists do information-processing analyses of complex skills, they find that the same kinds of basic problem-solving processes are used in task after task. Several examples of this have come up in the course of this essay. For example, although the original General Problem Solver (GPS), built to solve symbolic logic problems, was not in fact very general in the range of problems it could solve, the kinds of processes used by GPS appear over and over again in simulations of human performances on complex tasks. For example, means–ends analysis, generate-and-test routines, subgoal formation, and other kinds of planning are used in tasks as varied as inventing buggy arithmetic routines, planning compositions, constructing geometry proofs, and troubleshooting electronic devices. The reason that a single artificial intelligence program cannot solve a wide variety of problems is apparently not that the fundamental processes it applies are widely different across domains, but rather that the program must apply these processes to very specific, organized bodies of knowledge. Each simulation must build in the relevant knowledge, and so it becomes specific to its knowledge base (see Dehn & Schank, 1982).

Third, a variety of basic processes such as perceiving stimuli, encoding, classifying, generating responses, and executing responses are quite obviously involved in a number of different cognitive performances. These are the building block processes of intelligence and aptitude of the kind studied by Hunt and other students of the cognitive correlates of intelligence. Some years ago, Simon (1976), considering what information processing analyses of various tasks might suggest about the nature of intelligence, suggested that very low-level components (such as the building blocks) and very high-level ones (such as means–ends analysis and the like) are shared across many tasks and are therefore general abilities. The specific knowledge varies from task to task, however, producing the domain-specificity of cognitive abilities.

Finally, in some of the most recent and provocative work on the nature of intelligence, an apparently common body of "executive" or self-regulatory processes have been identified. Processes such as keeping track of one's own understanding or knowledge, initiating review or rehearsal activities when needed, deliberately organizing one's attention and other resources in order to learn something, or planning a set of actions so as to meet goals within the limits of certain constraints are all activities that have been shown to be characteristic of effective learners, good readers of texts, good writers, and strong problem solvers. These processes are relatively absent in younger or less intelligent individuals. These higher order or metacognitive skills, as they are often called, have become the object of an important recent line of research.

Self-Monitoring and Metacognition

Metacognition is surely one of the "boom" fields of recent cognitive psychology. The term *metacognition* is a relatively new one, whose field of reference has so exploded in just a few years that thoughtful scholars (cf., A. L. Brown, Bransford, Ferrara, & Campione, 1983) are beginning to suggest that it be abandoned as confusing, and that more specific terms such as *self-monitoring* and *self-regulation* should be substituted. The broad domain of metacognition includes (a) knowledge about cognition in general, (b) knowledge about one's own knowledge or cognitive strategies, and (c) application of these two kinds of knowledge to the planning and execution of appropriate mental activities in learning and problem-solving situations.

Several investigators have documented the fact that knowledge about cognition increases with age, and that older children are better able than younger ones to describe what one *ought* to do to remember something—for example, how to remember to take one's skates to school the next day. There is also a small body of evidence showing that younger children and the "developmentally young" (i.e., the re-

tarded) are less able than older children to assess when they have understood a message, when they are ready for a test, and the like. Convergent with this set of findings is a line of research showing that much of the deficit of retarded people in simple learning tasks, such as memorization, comes from their failure to apply well-known strategies such as rehearsal or information of mnemonics.

A number of studies have shown impressive gains in immediate performance on such tasks by simply instructing individuals to rehearse or to engage in verbal elaboration. In these studies, however, there was almost complete lack of transfer even to only slightly modified tasks. This led to a search for superordinate (Belmont, Butterfield, & Ferretti, 1982) or metacognitive skills (A. L. Brown, 1978), such as assessing one's own readiness for a test, apportioning study time, or deciding when to use rehearsal, imagery, or self-interrogation strategies that might promote general improvement. Some modest successes have been reported, but not enough for cognitive scientists to be convinced yet that even mild retardation can be overcome by training in superordinate skills of the kind studied thus far.

Most of the initial research on metacognitive training focused on memorization tasks that require very specialized kinds of strategies and that may have only a limited function outside of the laboratory and certain very specialized kinds of school learning (e.g., vocabulary lists). Recent research on memory, showing the power of chunked and organized knowledge in extending memory power (Chi, 1978), calls into question the extent to which strategies for artificial memorizing are likely to be an optimal approach to take even in simple school learning. A very recent shift toward the study of processes of self-control and self-regulation in more complex kinds of learning and performance, ranging from reading comprehension to writing compositions to learning new subject-matter domains, offers a more promising perspective on the development of learning abilities and the improvement of a variety of learning competencies.

Effects of reciprocal teaching. To illustrate this new perspective, I describe a recent training study which embodies many of the ideas under consideration in the field. The experiments were conducted by Palincsar and Brown (1984) with middle school children who had extremely weak reading comprehension skills. The children were divided into small groups, and with an adult, each group engaged in a process called "reciprocal teaching." The children took turns posing questions about and summarizing short texts that they read. The other members of the group commented on the quality of the questions or summaries and tried to help formulate better questions or summaries. Reciprocal teaching sessions were conducted daily for several weeks. During the intervention, there were daily assessments in which children individually read passages and answered questions about them. Assessments continued for several generalization days after the intervention ended,

and there also were 3 days of assessment after an 8-week break. In addition, there were some generalization and transfer tests. In the generalization tests, students read social studies and science texts in their regular classrooms and answered comprehension questions about them. The children were not aware that these tests had anything to do with the experimental teaching in which they had participated.

Table 1 contains some protocol segments from reciprocal teaching sessions with one of the weaker readers in the study. In these segments, there is a very great shift apparent in the child's ability to pose questions about the texts. At first the child cannot formulate questions at all and does not even produce full sentences. Later the child is able to pose coherent questions and by Day 15 is able to formulate a single question that addresses the main point of the passage. The protocol clearly shows the role of the adult in this process. At the beginning the teacher actually formulates the questions and the child does little more than repeat them. Later the adult provides portions of the questions, often the initial words of a sentence that will produce an appropriate question. By the end the child has taken over the entire process. Notice

Table 1
Protocol Excerpts Showing the Acquisition of Question-Asking by a Seventh-Grade Student (C) With Reciprocal Teaching

Day 1

TEXT: The water moccasin, somewhat longer than the copperhead, is found in the southeastern states. It lives in swampy regions. It belongs, as do also the copperhead and the rattlesnakes, to a group of poisonous snakes called pit vipers. They have pits between their eyes and their nostrils which, because they are sensitive to heat, help the snakes tell when they are near a warm-blooded animal. Another name for the water moccasin is "cottonmouth." This name comes from the white lining of the snake's mouth.

C: What is found in the southeastern snakes, also the copperhead, rattlesnakes, vipers—they have. I'm not doing this right.

T: All right. Do you want to know about the pit vipers?

C: Yeah.

T: What would be a good question about the pit vipers that starts with the word "why?"

C: (No response)

T: How about, "Why are the snakes called pit vipers?"

C: Why do they want to know that they are called pit vipers?

T: Try it again.

C: Why do they, pit vipers in a pit?

T: How about, "Why do they call the snakes pit vipers?"

C: Why do they call the snakes pit vipers?

T: There you go! Good for you.

Table 1, continued

Day 7	
TEXT:	Perhaps you are wondering where the lava and other volcanic products come from. Deep within our earth there are pockets of molten rock called magma. Forced upward in part by gas pressure, this molten rock continually tries to reach the surface. Eventually—by means of cracks in the crustal rocks or some similar zone of weakness—the magma may break out of the ground. It then flows from the vent as lava, or spews skyward as dense clouds of lava particles.
C:	How does the pressure from below push the mass of hot rock against the opening? Is that it?
T:	Not quite. Start your question with, "What happens when?"
C:	What happens when the pressure from below pushes the mass of hot rock against the opening?
T:	Good for you! Good job.
Day 15	
TEXT:	Scientists also come to the South Pole to study the strange lights that glow overhead during the Antarctic night. (It's a cold and lonely world for the few hardy people who "winter over" the polar night.) These "southern lights" are caused by the Earth acting like a magnet on electrical particles in the air. They are clues that may help us understand the Earth's core and the upper edges of its blanket of air.
C:	Why do scientists come to the south pole to study?
T:	Excellent question! That is what this paragraph is all about.

Note: T = teacher, C = student (Charles). Adapted from Palincsar and Brown (1984) by permission.

also that the standards for what the adult accepts as a good question from the child keep changing; stiffer requirements are applied toward the end than in the first days of the reciprocal teaching.

According to various measures this process greatly affected reading comprehension. Figure 9 shows increases in the percentage of assessment questions answered correctly by several students during the various phases of the first experiment. From a preintervention baseline measurement of very few questions answered correctly, most of the children moved up to about an 80 percent correct response rate, and they remained there after an 8-week break. Several control groups showed no increase in correct assessment responding during the same period. Figure 10 shows the effects of the training on the generalization tasks in regular classrooms. Again, the reciprocal teaching experience was shown to have a powerful effect on comprehension in a very different physical setting and under quite different measurement conditions. The experiment was replicated under more ordinary school conditions, as part of the regular instruction offered by the teacher in reading. Very similar results were obtained.

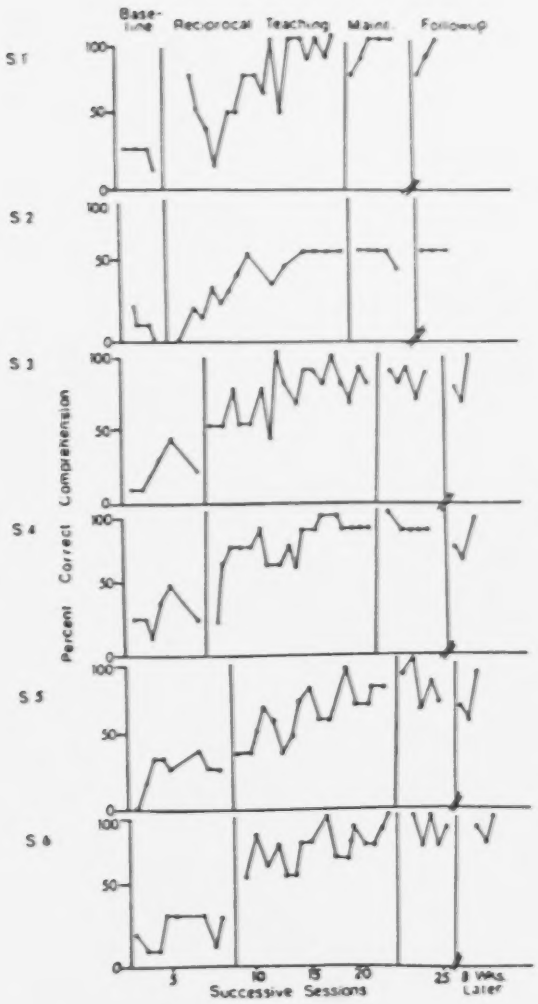

Figure 9. Results correct on the daily assessment passages for the six subjects in the Reciprocal Teaching Group of Study 1. Reprinted from Palincsar and Brown (1984) by permission.

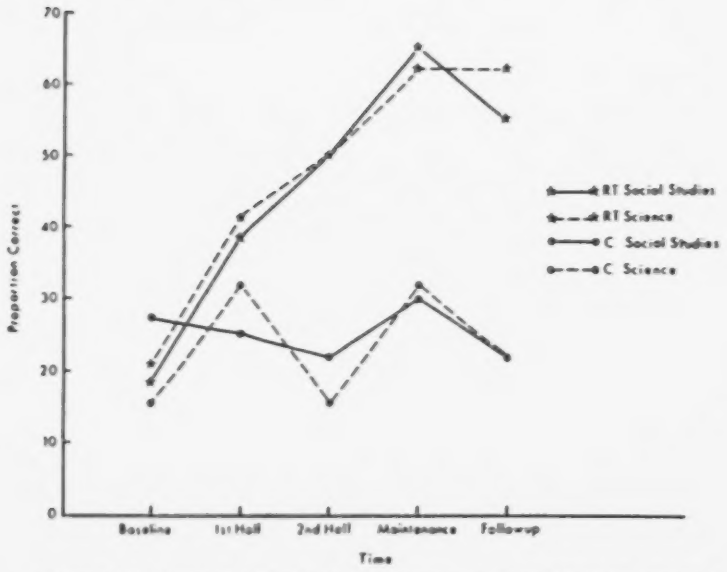

Figure 10. Percent correct on the classroom generalization probes of Study 1 for the Reciprocal Teaching (RT) and untreated control (C) groups in both science and social studies classes. Reprinted from Palincsar and Brown (1984) by permission.

The Palincsar and Brown study is not only a dramatic success story in educational intervention. It also illustrates many of the issues being studied in today's research on self-monitoring and cognitive strategy learning. First, the skills in question asking and summarizing that the children in the study practiced are probably not directly called upon in skilled reading. The automatic nature of many reading comprehension processes, the speed at which reading proceeds, and its sequential nature make it implausible that in the normal course of skilled reading, people actually pose questions or create summaries for themselves. There is, then, only an indirect relation between the strategies taught and what was probably involved in the children's subsequent relatively skilled reading performance. This indirect relation between strategies taught and skilled performance is also characteristic of findings in other metacognitive training research.

This raises the question of how instruction that focuses on overt, self-conscious strategies that are not components of skilled perfor-

mance might improve processes that progress quite automatically. The answer may lie in the fact that when readers use self-questioning and similar strategies, they evoke processes of inference and interpretation that eventually evolve into the automated information processing characteristic of skilled readers. This would mean that metacognitive strategies are better understood as aspects of cognitive skill acquisition than as aspects of cognitive skill expertise. With this assumption, researchers would be less interested in identifying components of skilled performance, and more interested in directly studying processes of learning. There is in fact a growing interest within cognitive science in the processes of cognitive acquisition (see Anderson, 1981) and some likelihood that this will become a dominant concern in the next decade, further linking developmental and experimental psychologists (cf. A. L. Brown et al., 1983).

Another possibility is that the strategies taught do not promote acquiring skilled processes so much as they activate or release capacities already available. The relative speed with which reading skill improved with reciprocal teaching suggests that this may be at least partially the case. As shown in Figure 10, correct answers to assessment questions increased very quickly after only a few sessions of reciprocal teaching for most of the students. Very rapid change in performance is also characteristic of the memory strategy training studies mentioned earlier. To the extent that strategy training releases rather than builds processing capability, one would expect this kind of instruction to be effective only if the relevant capabilities were already present.

There is some evidence that certain forms of metacognitive training can actually suppress performance, at least temporarily, if the knowledge or skills necessary to use new information is not already present. Scardamalia and Paris (in press) taught children to recognize and identify certain rhetorical devices that are known to affect the self-monitoring performance of skilled writers. This training increased the students' use of these rhetorical devices in their written compositions. However, no improvement, and even some depression, in overall organization and coherence of the compositions resulted. Scardamalia and Paris attributed this to a fundamental strategy that children use for composing, one they call "knowledge telling." The strategy involves little planning and the child writes down in sequence everything he or she thinks of relating to the topic. Because there was no overall planning, the children could not use these rhetorical devices to form the framework for a well-organized argument as adult writers would do.

A third important point is that reciprocal teaching is a special form of social interaction that may in fact be central to the acquisition of generalized cognitive skill. Traditional views of the way in which social interaction affects learning focus on the adult as provider of new information, as a modeler of correct performance, and as a selective rein-

forcer of children's tries at producing the performance. The reciprocal teaching of the Palincsar and Brown study was inspired by a different view of social processes in learning that is attracting increased attention among cognitive psychologists interested in the development of general cognitive competence. The Soviet psychologist Vygotsky (1978; see also Wertsch, 1978) has argued that cognition begins in social situations in which a child shares responsibility for producing a complete performance with an adult. The child does what he or she can, the adult the rest. In this way, practice on components occurs in the context of the full performance. In naturally occurring interactions of this kind, the adult will gradually increase expectations of how much of the full performance the child can be responsible for.

It should be clear that the Palincsar and Brown experiments should properly be regarded as more provocative than definitive. Their success in teaching a socially valuable skill, after many failures over the years, is stimulating. However, cognitive scientists do not really know what component in the reciprocal teaching method actually produced the success. So many elements of instruction were combined that it is impossible to determine from this study alone which parts of the instruction were essential. Further, it is not clear exactly what was taught. It is obvious that the children learned to ask questions and to summarize. However, the true target skill was neither of these, but rather skillful reading comprehension, and it is not completely clear why practice in asking questions and summarizing should produce that skill.

None of this is said in the spirit of criticizing the Palincsar and Brown experiments. Rather, it is said in order to emphasize that research on self-monitoring and metacognitive skills is at this time a highly promising but still largely unexplored domain. It is attracting considerable attention because there is some broad theory that suggests that it ought to work, and because a few studies such as the one cited have produced some dramatic successes. However, considerable caution in interpretation and in expectations for the future is necessary or psychologists risk another round of enthusiastic faddism.

Conclusion: Learning in the Future

The examples of cognitive research on learning and thinking developed in the course of this essay have been intended to convey the sense of excitement and of open possibilities that now pervade many branches of cognitive psychology. A major feature of current cognitive research is its focus on complex forms of intellectual competence. This has the effect of making large segments of fundamental research more immediately relevant to questions of instruction than has usually been the

case for psychological research on learning. As a result, a new cognitive instructional psychology is growing up as a special branch of cognitive psychology.

The emergence of instruction as an arena of concern for cognitive psychology is helping to focus the field's attention on questions of cognitive change. Instructional psychology seeks to formulate principles that can guide interventions designed to help people to learn. People learn, however, even when they are not taught, and so instruction must be construed as interventions in a learner's ongoing processes of knowledge acquisition. To develop principles of intervention, therefore, it is essential that we know what these acquisition processes are like. For this reason instructional psychology requires strong theories of the processes involved in cognitive change.

For a considerable period of time cognitive psychologists had given up the long-standing interest of experimental psychologists in questions of how changes in performance and competence come about. During this period, cognitive psychology focused instead on building detailed descriptions of given states of cognitive competence. Expert and novice states were often compared, but little was done to explain how people might pass from one state to the other. This inattention to processes of change is ending. In fact, the topic of learning is high on today's cognitive science agenda and will probably draw more and more attention in the next few years.

Most cognitive research on learning up to now has been concerned with accounting for changes in performance skills—that is, for developing speed and accuracy in doing things like solving algebra equations or programming computers. Some elegant and highly plausible theories of how early states of competence are transformed in the course of practice now exist. However, there has been little attention thus far to the question of how conceptual knowledge is acquired. Cognitive scientists can show how schemata influence learning from texts, for example, but their models of how the schemata themselves are learned are poorly developed. This lacuna in knowledge is widely recognized among cognitive scientists and some are now beginning to turn their attention to questions associated with the acquisition processes of conceptual learning. Psychologists can probably expect a new generation of cognitive learning theories to emerge in the next few years that will substantially modify the theoretical landscape. As this happens instructional questions are likely to become even more visible and central in cognitive psychology.

References

Anderson, J. R. (Ed.). (1981). *Cognitive skills and their acquisition.* Hillsdale, NJ: Erlbaum.

Anderson, R. C., Reynolds, R. E., Schallert, E. T., & Goetz, E. T. (1977). Frameworks for comprehending discourse. *American Educational Research Journal, 14,* 367–381.

Baroody, A. J., & Gannon, K. E. (1984). The development of the commutativity principle and economical addition strategies. *Cognition and Instruction, 1*(3), 321–339.

Bartlett, F. C. (1932). *Remembering: A study in experimental and social psychology.* New York: Cambridge University Press.

Belmont, J. M., Butterfield, E. C., & Ferretti, R. P. (1982). To secure transfer of training to instruct self-management skills. In D. K. Detterman & R. J. Sternberg (Eds.), *How and how much can intelligence be increased?* (pp. 147–155). Norwood, NJ: Ablex.

Bisanz, J., Danner, F., & Resnick, L. B. (1979). Changes with age in measures of processing efficiency. *Child Development, 50,* pp. 132–141.

Bransford, J. D., & Johnson, M. K. (1972). Contextual prerequisites for understanding: Some investigations of comprehension and recall. *Journal of Verbal Learning and Verbal Behavior, 11,* 717–726.

Briars, D. J., & Larkin, J. H. (1984). An integrated model of skill in solving elementary word problems. *Cognition and Instruction, 1*(3), 245–296.

Brown, A. L. (1978). Knowing when, where, and how to remember: A problem of metacognition. In R. Glaser (Ed.), *Advances in instructional psychology* (Vol. 1, pp. 77–165). Hillsdale, NJ: Erlbaum.

Brown, A. L., Bransford, J. D., Ferrara, R. A., & Campione, J. C. (1983). Learning, remembering, and understanding. In J. H. Flavell and E. M. Markman (Eds.), *Handbook of child psychology: Vol. 3. Cognitive development* (4th ed., pp. 420–494). New York: Wiley.

Brown, J. S., & Burton, R. R. (1978). Diagnostic models for procedural bugs in basic mathematical skills. *Cognitive Science, 2*(2), 155–192.

Brown, J. S., & VanLehn, K. (1982). Towards a generative theory of "bugs." In T. P. Carpenter, J. M. Moser, & T. A. Romberg (Eds.), *Addition and subtraction: A cognitive perspective* (pp. 117–135). Hillsdale, NJ: Erlbaum.

Burton, R. R. (1982). Diagnosing bugs in a simple procedural skill. In D. Sleeman & J. S. Brown (Eds.), *Intelligent tutoring systems* (pp. 157–183). New York: Academic Press.

Buswell, G. T. (1920). An experimental study of eye-voice span in reading. *Supplementary Educational Monographs,* No. 17. Chicago: University of Chicago Press.

Carpenter, T. P., Hiebert, J., & Moser, J. M. (1981). The effect of problem structure on first graders' initial solution processes for simple addition and subtraction problems. *Journal for Research in Mathematics Education, 12*(1), 27–39.

Carpenter, T. P., & Moser, J. M. (1982). The development of addition and subtraction problem-solving skills. In T. P. Carpenter, J. M. Moser, & T. A. Romberg (Eds.), *Addition and subtraction: A cognitive perspective* (pp. 9–24). Hillsdale, NJ: Erlbaum.

Carroll, J. B. (1976). Psychometric tests as cognitive tasks: A new "structure of intellect." In L. B. Resnick (Ed.), *The nature of intelligence* (pp. 27–56). Hillsdale, NJ: Erlbaum.

Carroll, J. B. (1982). The measurement of intelligence. In R. J. Sternberg (Ed.), *Handbook of human intelligence* (pp. 29–120). Cambridge, MA: Harvard University Press.

Carry, L. R., Lewis, C., & Bernard, J. E. (1980). *Psychology of equation solving—an information-processing study.* Austin: University of Texas.

Cattell, J. M. (1886). The time it takes to see and name objects. *Mind, 11,* 63–65.

Chi, M. T. H. (1978). Knowledge structures and memory development. In R. Siegler (Ed.), *Children's thinking: What develops?* (pp. 73–96). Hillsdale, NJ: Erlbaum.

Chi, M. T. H., Feltovich, P. J., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science, 5*(2), 121–152.

Chiesi, H. L., Spilich, G. J., & Voss, J. F. (1979). Acquisition of domain-related information in relation to high and low domain knowledge. *Journal of Verbal Learning and Verbal Behavior, 18,* 257–274.

Cooper, L. A. (1980). Spatial information processing: Strategies for research. In R. E. Snow, P. A. Federico, & W. E. Montague (Eds.), *Aptitude, learning, and instruction* (Vol. 1, pp. 149–176). Hillsdale, NJ: Erlbaum.

Cooper, L. A., & Regan, D. T. (1982). Attention, perception, and intelligence. In R. J. Sternberg (Ed.), *Handbook of human intelligence* (pp. 123–169). Cambridge, MA: Harvard University Press.

Cronbach, L. J. (1970). *Essentials of psychological testing* (3rd ed.). New York: Harper & Row.

Curtis, M. E., & Glaser, R. (1981). Changing conceptions of intelligence. *Review of Research in Education, 9,* 111–148.

Davis, R. B. (1983). Complex mathematical cognition. In H. P. Ginsburg (Ed.), *The development of mathematical thinking* (pp. 253–290). New York: Academic Press.

Dehn, N., & Schank, R. (1982). Artificial and human intelligence. In R. J. Sternberg (Ed.), *Handbook of human intelligence* (pp. 352–391). Cambridge, MA: Harvard University Press.

Ernst, G. W., & Newell, A. (1969). *GPS: A case study in generality and problem solving.* New York: Academic Press.

Ericsson, K. A., & Simon, H. A. (1984). *Protocol analysis: Verbal reports as data.* Cambridge, MA: MIT Press.

Flammer, A., & Kintsch, W. (Eds.). (1982). *Discourse processing: (Vol. 8) Advances in psychology.* Amsterdam: North-Holland.

Fleisher, L. S., & Jenkins, J. R. (1978). Effects of contextualized and decontextualized practice conditions on word recognition. *Learning Disabilities Quarterly, 1*(3), 39–47.

Frederiksen, C. H. (1979). Discourse comprehension and early reading. In L. B. Resnick & P. A. Weaver (Eds.), *Theory and practice of early reading* (Vol. 1, pp. 155–186). Hillsdale, NJ: Erlbaum.

Frederiksen, J. R., Warren, B. M., & Rosebery, A. S. (in press). A componential approach to training reading skills. *Cognition and Instruction.*

Friedman, M., Das, J., & O'Connor, N. (Eds.). (1980). *Intelligence and learning.* New York: Plenum Press.

Fuson, K. C. (1982). An analysis of the counting-on solution procedure in addition. In T. P. Carpenter, J. M. Moser, & T. A. Romberg (Eds.), *Addition and subtraction: A cognitive perspective* (pp. 67–81). Hillsdale, NJ: Erlbaum.

Gelman, R., & Gallistel, C. R. (1978). *The child's understanding of number.* Cambridge, MA: Harvard University Press.

Ginsburg, H. P. (1977). *Children's arithmetic: The learning process*. New York: Van Nostrand Reinhold.

Ginsburg, H. P. (Ed.). (1983). *The development of mathematical thinking*. New York: Academic Press.

Glaser, R. (1977). *Adaptive education: Individual diversity and learning*. New York: Holt, Rinehart & Winston.

Glaser, R. (1984). Education and thinking: The role of knowledge. *American Psychologist, 39*(1), 93–104.

Greeno, J. G. (1983, August). *Investigations of a cognitive skill*. Paper presented at the Annual meeting of the American Psychological Association, Anaheim, CA.

Greeno, J. G., Riley, M. S., & Gelman, R. (1984). Conceptual competence and children's counting. *Cognitive Psychology, 16*(1), 94–143.

Groen, G. J., & Parkman, J. M. (1972). A chronometric analysis of simple addition. *Psychological Review, 79*(4), 329–343.

Groen, G. J., & Poll, M. (1973). Subtraction and the solution of open sentence problems. *Journal of Experimental Child Psychology, 16*, 292–302.

Groen, G. J., & Resnick, L. B. (1977). Can preschool children invent addition algorithms? *Journal of Educational Psychology, 69*(6), 645–652.

Halliday, D., & Resnick, R. (1974). *Fundamentals of physics*. New York: Wiley.

Houlihan, D. M., & Ginsburg, H. P. (1981). The addition methods of first- and second-grade children. *Journal for Research in Mathematics Education, 12*(2), 95–106.

Huey, E. G. (1968). *The psychology and pedagogy of reading*. Cambridge, MA: MIT Press. (Original work published 1908)

Hunt, E. (1978). Mechanics of verbal ability. *Psychological Review, 85*, 109–130.

Hunt, E., Frost, N., & Lunneborg, C. (1973). Individual differences in cognition: A new approach to intelligence. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 7, pp. 87–122). New York: Academic Press.

Kieras, D. (1977). Problems of reference in text comprehension. In M. A. Just & P. A. Carpenter (Eds.), *Cognitive processes in comprehension* (pp. 249–270). Hillsdale, NJ: Erlbaum.

Kintsch, W. (1979). On modeling comprehension. *Educational Psychologist, 14*, 3–14.

Kintsch, W., & vanDijk, T. A. (1978). Toward a model of text comprehension and production. *Psychological Review, 85*, 363–394.

Lesgold, A. M., & Resnick, L. B. (1983). How reading difficulties develop: Perspectives from a longitudinal study. In J. P. Das, R. Mulcahy, & A. E. Wall (Eds.), *Theory and research in learning disability* (pp. 155–188). New York: Plenum Press.

Magone, M. E. (1977). *Goal analysis and feature detection as processes in the solution of an insight problem*. Unpublished master's thesis, University of Pittsburgh, Pittsburgh, PA.

Mandler, J. M. (1978). A code in the node: The use of a story schema in retrieval. *Discourse Processes, 1*, 14–35.

Mandler, J. M., & Johnson, N. S. (1977). Remembrance of things parsed: Story structure and recall. *Cognitive Psychology, 9*, 111–151.

Matz, M. (1982). Towards a process model for high school algebra errors. In D. Sleeman & J. S. Brown (Eds.), *Intelligent tutoring systems* (pp. 25–50). New York: Academic Press.

Meyer, B. J. F. (1975). *The organization of prose and its effect on recall.* Amsterdam: Elsevier.

Meyer, B. J. F. (1984). Text dimensions and cognitive processing. In H. Mandl, N. L. Stein, & T. Trabasso (Eds.), *Learning and comprehension of text* (pp. 3–51). Hillsdale, NJ: Erlbaum.

Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review, 63,* 81–97.

Mulholland, T. M., Pellegrino, J. W., & Glaser, R. (1980). Components of geometric analogy solution. *Cognitive Psychology, 12,* 252–284.

Neches, R. (1981). *Models of heuristic procedure modification.* Unpublished doctoral dissertation, Carnegie-Mellon University, Pittsburgh, PA.

Nesher, P. (1982). Levels of description in the analysis of addition and subtraction word problems. In T. P. Carpenter, J. M. Moser, & T. A. Romberg (Eds.), *Addition and subtraction: A cognitive perspective* (pp. 25–38). Hillsdale, NJ: Erlbaum.

Nesher, P., & Teubal, E. (1975). Verbal cues as an interfering factor in verbal problem solving. *Educational Studies in Mathematics, 6,* 41–51.

Newell, A., & Simon, H. A. (1972). *Human problem solving.* Englewood Cliffs, NJ: Prentice-Hall.

Nickerson, R. S., Salter, W., Shepard, S., & Herrnstein, J. (1984). *The teaching of learning strategies.* Cambridge, MA: Bolt, Beranek, and Newman.

Palincsar, A. S., & Brown, A. L. (1984). Reciprocal teaching of comprehension-fostering and comprehension-monitoring activities. *Cognition and Instruction, 1*(2), 117–175.

Papert, S. A. (1980). *Mindstorms.* New York: Basic Books.

Pellegrino, J. W., & Glaser, R. (1979). Cognitive correlates and components in the analysis of individual differences. *Intelligence, 3,* 187–214.

Pellegrino, J. W., & Glaser, R. (1980). Components of inductive reasoning. In R. E. Snow, P. A. Federico, & W. E. Montague (Eds.), *Aptitude, learning and instruction* (Vol. 1, pp. 177–218). Hillsdale, NJ: Erlbaum.

Perfetti, C. A. (Ed.). (in press). *Reading Ability.* New York: Oxford University Press.

Piaget, J. (1965). *The child's conception of number.* New York: Norton. (Original work published 1941)

Resnick, L. B. (Ed.). (1976). *The nature of intelligence.* Hillsdale, NJ: Erlbaum.

Resnick, L. B. (1981a). Instructional psychology. In M. R. Rosenzweig & L. W. Porter (Eds.), *Annual review of psychology* (Vol. 32, pp. 659–704).

Resnick, L. B. (1981b). Social assumptions as a conflict for science: Some reflections on psychology and education. *Educational Psychologist, 16*(1), 1–10. [University of Pittsburgh, Learning Research and Development Center, LRDC Reprint Series 1981/18.]

Resnick, L. B. (1982). Syntax and semantics in learning to subtract. In T. P. Carpenter, J. M. Moser, & T. A. Romberg (Eds.), *Addition and subtraction: A cognitive perspective* (pp. 136–155). Hillsdale, NJ: Erlbaum.

Resnick, L. B. (1983). A developmental theory of number understanding. In H. P. Ginsburg (Ed.), *The development of mathematical thinking* (pp. 109–151). New York: Academic Press.

Resnick, L. B., & Neches, R. (1984). Factors affecting individual differences in learning ability. In R. J. Sternberg (Ed.), *Advances in the psychology of human intelligence* (Vol. 2, pp. 275–323) Hillsdale, NJ: Erlbaum.

Resnick, L. B., & Omanson, S. F. (in press). Learning to understand arithmetic. In R. Glaser (Ed.), *Advances in instructional psychology* (Vol. 3). Hillsdale, NJ: Erlbaum.

Riley, M. S., Greeno, J. G., & Heller, J. I. (1983). The development of children's problem-solving ability in arithmetic. In H. P. Ginsburg (Ed.), *The development of mathematical thinking* (pp. 153–196). New York: Academic Press.

Rumelhart, D. E., & McClelland, J. L. (1981). Interactive processing through spreading activation. In A. M. Lesgold & C. A. Perfetti (Eds.), *Interactive processes in reading* (pp. 37–60). Hillsdale, NJ: Erlbaum.

Scardamalia, M., & Paris, P. (in press). The function of explicit discourse knowledge in the development of text representations. *Cognition and Instruction.*

Schank, R., & Abelson, R. (1977). *Scripts, plans, goals, and understanding: An inquiry into human knowledge structures.* Hillsdale, NJ: Erlbaum.

Segal, J. W., Chipman, S. F., & Glaser, R. (Eds.). (1985). *Thinking and learning skills: Relating instruction to research,* (Vol. 1). Hillsdale, NJ. Erlbaum.

Simon, H. A. (1976). Identifying basic abilities underlying intelligent performance of complex tasks. In L. B. Resnick (Ed.), *The nature of intelligence* (pp. 65–98). Hillsdale, NJ: Erlbaum.

Sleeman, D. (1982). Assessing aspects of competence in basic algebra. In D. Sleeman & J. S. Brown (Eds.), *Intelligent tutoring systems* (pp. 185–199). New York: Academic Press.

Snow, R. E. (1976). Research on aptitudes: A progress report. In L. S. Shulman (Ed.), *Review of research in education* (Vol. 4, pp. 50–105). Itasca, IL: Peacock.

Snow, R. E., Federico, P. A., & Montague, W. E. (Eds.). (1980). *Aptitude, learning, and instruction* (Vols. 1–2). Hillsdale, NJ: Erlbaum.

Spilich, G. J., Vesonder, G. T., Chiesi, H. L., & Voss, J. F. (1979). Text-processing of domain-related information for individuals with high and low domain knowledge. *Journal of Verbal Learning and Verbal Behavior, 18,* 275–290.

Steffe, L., Thompson, P., & Richards, J. (1982). Children's counting in arithmetical problem solving. In T. P. Carpenter, J. M. Moser, & T. A. Romberg (Eds.), *Addition and subtraction: A cognitive perspective* (pp. 83–97). Hillsdale, NJ: Erlbaum.

Stein, N. L., & Glenn, C. G. (1979). An analysis of story comprehension in elementary school children. In R. O. Freedle (Ed.), *New directions in discourse processing* (Vol. 2, pp. 53–120). Norwood, NJ: Ablex.

Stein, N. L., & Trabasso, T. (1983). What's in a story: An approach to comprehension and instruction. In R. Glaser (Ed.), *Advances in instructional psychology* (Vol. 2, pp. 213–267). Hillsdale, NJ: Erlbaum.

Sternberg, R. J. (1977a). Component processes in analogical reasoning. *Psychological Review, 84*(4), 353–378.

Sternberg, R. J. (1977b). *Intelligence, information processing, and analogical reasoning: The componential analysis of human abilities.* Hillsdale, NJ: Erlbaum.

Sternberg, R. J. (1980). Sketch of a componential subtheory of human intelligence. *Behavior and Brain Sciences, 3,* 573–584.

Sternberg, R. J., & Detterman, D. K. (Eds.). (1979). *Human intelligence: Perspectives on its theory and measurement.* Norwood, NJ: Ablex.

Svenson, O., & Broquist, S. (1975). Strategies for solving simple addition problems: A comparison of normal and subnormal children. *Scandinavian Journal of Psychology, 16,* 143–151.

Svenson, O., & Hedenborg, M. L. (1979). Strategies used by children when solving simple subtractions. *Acta Psychologica, 43*, 1–13.

Svenson, O., Hedenborg, M. L., & Lingman, L. (1976). On children's heuristics for solving simple additions. *Scandinavian Journal of Educational Research, 20*, 161–173.

Thibadeau, R., Just, M. A., & Carpenter, P. A. (in press). A model of the time course and content of reading. *Cognitive Science*.

Thorndike, E. L. (1922). *The psychology of arithmetic*. New York: Macmillan.

Thorndyke, P. W. (1977). Cognitive structures in comprehension and memory of narrative discourse. *Cognitive Psychology, 9*(1), 77–110.

Trabasso, T., Secco, T., & van den Broek, P. (1984). Causal cohesion and story coherence. In H. Mandl, N. L. Stein, & T. Trabasso (Eds.), *Learning and comprehension of text* (pp. 83–111). Hillsdale, NJ: Erlbaum.

Venezky, R. L., & Massaro, D. W. (1979). The role of orthographic regularity in word recognition. In L. B. Resnick & P. A. Weaver (Eds.), *Theory and practice of early reading*, (Vol. 1, pp. 85–107). Hillsdale, NJ: Erlbaum.

Vergnaud, G. (1982). A classification of cognitive tasks and operations of thought involved in addition and subtraction problems. In T. P. Carpenter, J. M. Moser, & T. A. Romberg (Eds.), *Addition and subtraction: A cognitive perspective* (pp. 39–59). Hillsdale, NJ: Erlbaum.

Voss, J. F. (1984). On learning and learning from text. In H. Mandl, N. L. Stein, & T. Trabasso (Eds.), *Learning and comprehension of text* (pp. 193–212). Hillsdale, NJ: Erlbaum.

Vygotsky, L. S. (1978). *Mind in society: The development of higher psychological process*. M. Cole, V. John-Steiner, S. Scribner, & E. Souberman (Eds. and Trans.). Cambridge, MA: Harvard University Press.

Wertsch, J. V. (1978). Adult-child interaction and the roots of metacognition. *Quarterly Newsletter of the Institute for Comparative Human Development, 1*, 15–18.

Woods, S. S., Resnick, L. B., & Groen, G. J. (1975). An experimental test of five process models for subtraction. *Journal of Educational Psychology, 67*(1), 17–21.

The Master Lecture Series, a continuing education program of five distinguished lectures, has been presented annually at the American Psychological Association convention since 1974. The lectures provide updates on the relevant issues, methodologies, and advances in research in a different area of specialization each year. This volume, *Psychology and Learning*, is the fourth in a series of bound volumes derived from the Master Lecture presentations. The volume should appeal to scholars and professional practitioners, to students who are just beginning their training in the field, and to nonpsychologists who wish to stay informed about current issues in the psychology of learning.

American Psychological Association
1200 Seventeenth Street, N.W.
Washington, D.C. 20036